

THE STRUCTURE THEORY OF NILSPACES I

YONATAN GUTMAN, FREDDIE MANNERS AND PÉTER P. VARJÚ

ABSTRACT. This paper forms the first part of a series by the authors [GMV16a, GMV16b] concerning the structure theory of *nilspaces* of Antolín Camarena and Szegedy. A nilspace is a compact space X together with closed collections of *cubes* $C^n(X) \subseteq X^{2^n}$, $n = 1, 2, \dots$ satisfying some natural axioms. Antolín Camarena and Szegedy proved that from these axioms it follows that (certain) nilspaces are isomorphic (in a strong sense) to an inverse limit of nilmanifolds. The aim of our project is to provide a new self-contained treatment of this theory and give new applications to topological dynamics.

This paper provides an introduction to the project from the point of view of applications to higher order Fourier analysis. We define and explain the basic definitions and constructions related to cubespaces and nilspaces and develop the *weak structure theory*, which is the first stage of the proof of the main structure theorem for nilspaces. Vaguely speaking, this asserts that a nilspace can be built as a finite tower of extensions where each of the successive fibers is a compact abelian group.

We also make some modest innovations and extensions to this theory. In particular, we consider a class of maps that we term *fibrations*, which are essentially equivalent to what are termed *fiber-surjective morphisms* by Antolín Camarena and Szegedy; and we formulate and prove a relative analogue of the weak structure theory alluded to above for these maps. These results find applications elsewhere in the project.

CONTENTS

1. An introduction to the project	2
2. Nilmanifolds and their Host–Kra cubes	7
3. Cubespaces and nilspaces	11
4. The structure theorem	17
5. A more detailed summary of the argument	19
6. The weak structure theory	27
7. The relative weak structure theory	36
Appendix A. An extended exposition of Host–Kra cubes	50
References	61

Date: May 31, 2016.

Key words and phrases. Gowers norms, higher order Fourier analysis, cubespace, nilspace, nilmanifold, Lie groups, fibration.

YG was supported by the ERC Grant *Approximate Algebraic Structures and Applications*. PPV was supported by the Royal Society.

1. AN INTRODUCTION TO THE PROJECT

This is the first in a series of three papers by the authors concerned with the structure theory of *cubespaces*, the others being [GMV16a] and [GMV16b]. Informally, a cubespace is a compact metric space X , together with some notion of when a collection of 2^k points of X form a “ k -cube”, subject to certain further axioms.

The study of cubespaces as axiomatic objects is becoming established as a major theme in the nascent area of higher order Fourier analysis. This programme has its origins in work of Host and Kra [HK08], where these objects appeared under the name of “parallelepiped structures”. The study of these objects was furthered by Antolín Camarena and Szegedy [ACS12], who in the same work formulated a strong structure theorem for cubespaces, subject to certain further hypotheses. The gist of the structure theorem is that, subject to these extra assumptions, all cubespaces arise in some sense from *nilmanifolds* $X = G/\Gamma$; i.e. they come from compact homogeneous spaces of nilpotent Lie groups.

The papers of Candela [Can16b, Can16a] expand on [ACS12], providing more detailed proofs. He also includes several additional results implicit in [ACS12], particularly about continuous systems of measures.

This structure theory has applications in two broad areas: additive combinatorics, and in particular the inverse theorem for the Gowers norms; and topological dynamics and ergodic theory. In this paper, we will approach this project from the point of view of someone interested primarily in understanding the former, and in particular Szegedy’s proof [Sze12] of the inverse theorem for the Gowers norms, which relies heavily on this structural result. The third paper in the series [GMV16b] will contain an introduction to the project focussed instead on applications to topological dynamics. Although the results proved in each paper are strongly relevant to the other, a reader more interested in the dynamical perspective might prefer to start there and refer to this paper only later.

The main purpose of our entire project is to provide a self-contained proof of the main structural result of [ACS12]. In many places our approach will follow that of [ACS12], or of previous work [HK05, HM07, HK08, HKM10, GT10], very closely. At other times, we give notably different arguments: sometimes because they are arguably simpler; sometimes to avoid certain technical difficulties (although perhaps at the expense of introducing different ones); and sometimes to obtain sharper conclusions. Furthermore, we do obtain new results in particular in the dynamical setting, and some of our proofs are optimized so as to prove these concurrently.

Our primary goal, however, is to obtain a fuller understanding of cubespaces and related structures. In our view, this understanding will continue to find new relevance as the field of higher order Fourier analysis matures.

1.1. Obstructions to Gowers uniformity. It is not immediately clear why the structure theory of cubespaces should be relevant to proving an inverse theorem for the Gowers norms: indeed, the deduction of the latter from the former is by no means straightforward [Sze12]. To provide some motivation, we will instead sketch a proof of a kind of converse. That is, we will argue that cubespaces

(subject to certain additional hypotheses) are obstructions to Gowers uniformity, and that therefore the inverse theorem itself implies that cubespaces are somehow related to nilmanifolds.

We will not recall in full all the relevant definitions (of the uniformity norms, nilmanifolds, polynomial maps, nilsequences etc.), referring the reader to [GT10, Gre15] but will informally sketch the set-up to motivate our discussion. For notational simplicity we will focus only on the U^3 norm, although these remarks apply more generally.


- (i) Given a function $f: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$ for N a prime (say), the uniformity norm $\|f\|_{U^3}$ is defined in terms of an average over *cube* or *parallelepiped* configurations in $\mathbb{Z}/N\mathbb{Z}$, e.g.

$$\|f\|_{U^3}^8 = \mathbb{E}_{c \in C^3(\mathbb{Z}/N\mathbb{Z})} f(c_{000}) \overline{f(c_{001})} \overline{f(c_{010})} \overline{f(c_{011})} \overline{f(c_{100})} \overline{f(c_{101})} \overline{f(c_{110})} \overline{f(c_{111})}$$

where $C^3(\mathbb{Z}/N\mathbb{Z}) \subseteq (\mathbb{Z}/N\mathbb{Z})^8$ consists of all tuples

$$(c_\omega)_{\omega \in \{0,1\}^3} = x + \omega_1 h_1 + \omega_2 h_2 + \omega_3 h_3$$

for $x, h_1, h_2, h_3 \in \mathbb{Z}/N\mathbb{Z}$. It is informative to think of these tuples as elements written on the vertices of a 3-dimensional cube, as shown.



- (ii) Given a nilmanifold G/Γ (with suitable additional structure) there is also a notion of cubes on G/Γ , given by a construction due to Host and Kra [HK05, HK08]. Specifically, suppose G is a 2-step nilpotent Lie group and Γ a discrete co-compact subgroup; then there is a closed subset $C^3(G/\Gamma) \subseteq (G/\Gamma)^8$ somewhat analogous to the parallelepipeds in an abelian group.
- (iii) There is a plentiful supply of maps $p: \mathbb{Z}/N\mathbb{Z} \rightarrow G/\Gamma$ which send cubes to cubes; that is, $p(c) \in C^3(G/\Gamma)$ for any $c \in C^3(\mathbb{Z}/N\mathbb{Z})$ (with p applied pointwise).
- (iv) The cubes on G/Γ satisfy a *corner constraint*: given $c \in C^3(G/\Gamma)$, if we know c_{000}, \dots, c_{110} then the last vertex c_{111} is uniquely determined by the others.
- (v) By a *nilsequence* on $\mathbb{Z}/N\mathbb{Z}$ we mean a function of the form $\phi = F \circ p$ where p is as above and $F: G/\Gamma \rightarrow \mathbb{C}$ is Lipschitz.

The inverse theorem states roughly that (for very large N) if $f: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$, $|f| \leq 1$ has $\|f\|_{U^3} \geq \delta$ then $|\langle f, \phi \rangle| \gtrsim_\delta 1$ for some nilsequence ϕ whose “complexity” is bounded in terms of δ . One can say that nilsequences are the only “obstructions to Gowers uniformity”: the only reason for a function to have large Gowers norm is if it correlates with a nilsequence.

All known proofs of this statement are significantly hard. By contrast, the converse statement – if a function correlates with a nilsequence then it has large Gowers norm – is quite straightforward. We will sketch a result along these lines now, although the details of the proof are not important in what follows.¹

Claim. *Suppose $\phi = F \circ p$ is a nilsequence in the sense of (v) with $\|F\|_\infty \leq 1$. Let $f: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$ be such that $|f| \leq 1$ and $|\langle f, \phi \rangle| \geq \delta$. Then $\|f\|_{U^3} = \Omega_{F,\delta}(1)$.*

The key point is that the lower bound on $\|f\|_{U^3}$ depends only on the choice of F (and so implicitly of G/Γ) and on δ ; *not* on N or p .

Proof sketch. By property (iv), there is a closed subset $Y \subseteq (G/\Gamma)^7$ and a function $\tau: Y \rightarrow (G/\Gamma)$ such that

$$C^3(G/\Gamma) = \{(\tau(y), y) : y \in Y\}.$$

Hence we get a continuous function $F \circ \tau$ on Y . By Tietze extension theorem, we can extend this to a bounded continuous function H on $(G/\Gamma)^7$. Any continuous function on a product space can be approximated (up to a small error in L^∞) by a finite sum of products of functions on the factors: that is, we can decompose

$$H(x_1, \dots, x_7) = \sum_{i=1}^k R_1^{(i)}(x_1) \dots R_7^{(i)}(x_7) + H_{\text{err}}$$

for some bounded continuous functions $R_j^{(i)}: G/\Gamma \rightarrow \mathbb{C}$, and some bounded continuous $H_{\text{err}}: (G/\Gamma)^7 \rightarrow \mathbb{C}$ such that $\|H_{\text{err}}\|_\infty = o_{F;k \rightarrow \infty}(1)$.

Now, for any x, h_1, h_2, h_3 in $\mathbb{Z}/N\mathbb{Z}$ we have that $(p(x), p(x+h_1), p(x+h_2), p(x+h_1+h_2), \dots)$ is in $C^3(G/\Gamma)$, and so

$$\begin{aligned} \phi(x) &= F(p(x)) \\ &= F(\tau(p(x+h_1), p(x+h_2), p(x+h_1+h_2), \dots, p(x+h_1+h_2+h_3))) \\ &= \sum_{i=1}^k R_1^{(i)}(p(x+h_1)) R_2^{(i)}(p(x+h_2)) \dots R_7^{(i)}(p(x+h_1+h_2+h_3)) + o_{k \rightarrow \infty}(1). \end{aligned}$$

Since $|\langle f, \phi \rangle|$ is bounded away from zero, we deduce that

$$\left| \mathbb{E}_{x, h_1, h_2, h_3} f(x) \overline{R_1^{(i)}(p(x+h_1)) R_2^{(i)}(p(x+h_2)) \dots R_7^{(i)}(p(x+h_1+h_2+h_3))} \right|$$

is bounded away from zero for some $i \in \{1, \dots, k\}$ (after choosing k appropriately in terms on F and δ). But this expression is a “Gowers inner product” of eight functions, and by the Gowers–Cauchy–Schwarz inequality (essentially multiple applications of Cauchy–Schwarz), this quantity is bounded above by

$$\|f\|_{U^3} \left\| R_1^{(i)} \circ p \right\|_{U^3} \dots \left\| R_7^{(i)} \circ p \right\|_{U^3}$$

¹This argument appears in [GT08, Proposition 12.6].

and noting that $\|R_j^{(i)} \circ p\|_{U^3} \leq \|R_j^{(i)}\|_\infty$ which is bounded, we get a lower bound on $\|f\|_{U^3}$ as required. \square

The key point is that the only properties of nilmanifolds, nilsequences etc. that we have used are those described in (ii), (iii) and (iv) above. So we have in fact shown:

Observation 1.1. *If X is any compact metric space equipped with some suitable notion of “cubes” as in (ii), having an abundance of cube-preserving maps $\mathbb{Z}/N\mathbb{Z} \rightarrow X$ as in (iii), and satisfying a corner constraint as in (iv), then functions of the form $F \circ p$ as in (v) obstruct Gowers uniformity on $\mathbb{Z}/N\mathbb{Z}$ in the sense of the above claim.*

Let us refer to such a space informally for now as a “nil-object” (the formal notion of a *nilspace* will be introduced later). Then the above observation can be summarized as follows.

Slogan 1.2. *Any “nil-object” is an obstruction to Gowers uniformity.*

But now, the inverse theorem for the U^3 norm tells us that any function $F \circ p$ coming from this construction must have something to do with a genuine nilsequence. Given some technical hypotheses, one can deduce that any such nil-object X must be very closely related to an actual nilmanifold G/Γ .

The conclusion of work of Szegedy [Sze12] is that it is possible to go in the other direction. He argues that all functions f with $\|f\|_{U^3}$ somewhat large correlate with something of the form $F \circ p$ where $p: \mathbb{Z}/N\mathbb{Z} \rightarrow X$ and $F: X \rightarrow \mathbb{C}$ is continuous, for *some* space X equipped with a notion of cubes, and some cube-preserving p , obeying some fairly reasonable additional axioms. Moreover, this proof is essentially purely analytic in nature, making no attempt to say anything about the structure of X .

Assuming this, we conclude:

Slogan 1.3. *The class of all “nil-objects” corresponds precisely to the obstructions to Gowers uniformity.*

Hence, the inverse theorem for the Gowers norms is essentially *equivalent* to classifying nil-objects X , showing in effect that they are all – essentially – nilmanifolds. This structural result is the goal of [ACS12].

1.2. An outline of this paper. The formal notion capturing properties (ii)-(iv) above and replacing the informal concept of a “nil-object”, is what we term a *nilspace*. A nilspace is a compact topological space satisfying certain “nilspace axioms”, which are both very abstract (e.g. they do not explicitly mention any group structure) but simultaneously strong enough for a strong structural result to hold.

The remaining tasks of this paper are as follows:

- (i) to explain the nilspace axioms formally;
- (ii) to state a precise version of the structure theorem;
- (iii) to outline the stages in the proof of this theorem; and

(iv) to prove some weaker structural results that constitute the first stage of the proof.

The remaining parts of the proof of the structure theorem appear in the companion papers [GMV16a, GMV16b], and we will provide pointers to the relevant sections of these works in the outline.

Most of our discussion will be of an expository or heuristic nature. The only parts that constitute steps in the rigorous proof of the structure theorem are Section 7, and in a way the definitions in Section 3. A reader already familiar with this whole approach and seeking only a complete proof could therefore read only these, together with the companion papers [GMV16a, GMV16b]; needless to say we do not recommend this strategy.

Most results proved in this paper are due to Antolín Camarena and Szegedy [ACS12], and the majority of the ideas we will discuss originate in [HK08] or [ACS12]. The primary exception is the “relative” version of this “weak structure theory” treated in Section 7. This generalization is fairly mechanical when stated in the language of *fibrations*, which are a class of maps between cubespaces satisfying certain additional hypotheses. This notion is almost equivalent to that of a *fiber-surjective morphism* appearing in [ACS12]; indeed, the latter is only defined for maps between two nilspaces, and in that case the definitions are equivalent. Our reasons for introducing the former are twofold: the definition of a fibration makes sense in greater generality, which we do actually use; and in our view, the alternative definition makes the analogy between relative and non-relative versions more transparent.

Before embarking on (i), we recall a small amount of background about Host–Kra cube groups, which puts the abstract definitions in some context. This is done in Section 2.

The definitions themselves are then expounded in Section 3.

With these in place, we are in a position to state the full structural result in Section 4. Then, we give a very heuristic outline of the high-level stages of the proof.

This outline is expanded upon in Section 5, which gives a detailed overview of each stage of the proof.

The final two sections, Section 6 and Section 7, are concerned with the first of these stages, which we have been calling the “weak structure theory”. Section 6 deals with the “standard” theory as it appears in [ACS12]. In Section 7, we consider *relative* versions of essentially all the elementary theory, which generalize the earlier results and which will be needed elsewhere in this project.

In the Appendix, we resume the discussion of Host–Kra cubes from Section 2, and make a detailed study of cubes on nilmanifolds in some example cases.

1.3. Acknowledgments. First and foremost we owe gratitude to Bernard Host who introduced us to the subject and to Omar Antolín Camarena and Balázs Szegedy whose groundbreaking work [ACS12] was a constant inspiration for us.

We would like to thank Emmanuel Breuillard, Jérôme Buzzi, Yves de Cornulier, Sylvain Crovisier, Eli Glasner, Ben Green, Bernard Host, Michał Rams, Balázs Szegedy, Anatoly Vershik and Benjamin

Weiss for helpful discussions. We are grateful to Pablo Candela and Bryna Kra for a careful reading of a preliminary version.

2. NILMANIFOLDS AND THEIR HOST–KRA CUBES

The key motivating examples of nilspaces (being in some sense the only examples) are nilmanifolds G/Γ equipped with their Host–Kra cube structures. Certainly anything we define or prove about general nilspaces should hold true for these spaces, and they provide a good source of intuition to guide the definitions and arguments in the abstract setting.

With this in mind, we will now briefly recall the relevant constructions and a few properties. A much more substantial exposition of this theory, including a number of examples explored in depth, is given in Appendix A; this also includes omitted proofs from this section. The reader unfamiliar with this area might wish to read this exposition before continuing with the bulk of the paper.

The notions and ideas presented in this section originate from [HK05, HK08, GT10].

We first recall the notion of a filtration.

Definition 2.1. Let G be a topological group. By a *filtration of degree s* on G , we mean a sequence $G = G_0 \supseteq G_1 \supseteq \cdots \supseteq G_{s+1} = \{\text{id}\}$ of closed subgroups, with the property that $[G_i, G_j] \subseteq G_{i+j}$ for all $i, j \geq 0$. (Here $[\cdot, \cdot]$ denotes commutation, and we use the convention $G_i = \{\text{id}\}$ for all $i \geq s+1$.)

The filtration is called *proper* if $G_0 = G_1$.


For most of the theory we will only consider proper filtrations, but we find it useful to permit $G_0 \neq G_1$ in the definition. Note that if G admits a proper filtration of degree s , then G is necessarily nilpotent, with nilpotency class at most s .

We write G_\bullet in place of G if we wish to emphasize that a group is equipped with a particular filtration.

The standard example of a filtration is the *lower central series* filtration on a group, given by $G_0 = G_1 = G$ and $G_{i+1} = [G, G_i]$ for each $i \geq 1$ (for a proof that this is indeed a filtration see [MKS66, Theorem 5.3]). This is in some well-defined sense the minimal proper filtration on G .


The fundamental construction concerning filtered groups is that of the *Host–Kra cube groups*. These are designed to be the appropriate analogues of parallelepipeds in abelian groups, in the setting of general filtered groups G_\bullet . To describe them, we will first set up some notation.

For a set X , we use the notation $X^{\{0,1\}^k}$ to mean the space of all functions $\{0,1\}^k \rightarrow X$. Concretely, this is just X^{2^k} but the reader should always imagine the elements written at the vertices of the discrete cube $\{0,1\}^k$, as in:



We denote by $[k]$ the set $\{1, \dots, k\}$. We find it convenient to identify subsets of $[k]$ with vertices of the discrete cube $\{0, 1\}^k$, mapping sets to their indicator functions. In particular, we write $\omega_1 \subseteq \omega_2$ for two vertices if $\omega_1(j) \leq \omega_2(j)$ for all $j \in [k]$.

By a *face* of the discrete cube $\{0, 1\}^k$ we mean a sub-cube obtained by fixing some subset of the coordinates. An *upper face* is one obtained by fixing some subset of the coordinates to equal 1; we can write this as $\{\omega \in \{0, 1\}^k : \omega \supseteq S\}$ for some $S \subseteq [k]$. For instance, in



an upper face of codimension two and a face of codimension one are indicated.


Definition 2.2. Let G_\bullet be a filtered topological group, and let $k \geq 0$ be an integer. The k -th Host–Kra cube group, denoted $\text{HK}^k(G_\bullet)$, is a sub-group of $G^{\{0,1\}^k}$ defined as follows.

For a subset $S \subseteq [k]$, let $F_S \subseteq \{0, 1\}^k$ denote the upper face corresponding to S , i.e. $F_S = \{\omega \in \{0, 1\}^k : \omega \supseteq S\}$. For any $x \in G$ and any face F of $\{0, 1\}^k$, let $[x]_F$ denote the configuration $\{0, 1\}^k \rightarrow G$ given by

$$[x]_F(\omega) = \begin{cases} x & : \omega \in F, \\ \text{id} & : \text{otherwise} . \end{cases}$$


Then $\text{HK}^k(G_\bullet)$ is the group generated by elements of the form $[x]_{F_S}$ where $S \subseteq \{0, 1\}^n$ and $x \in G_{|S|} = G_{\text{codim}(F_S)}$.

We will illustrate this in the case $k = 3$. First, taking $S = \emptyset$ we are free to take any constant configuration:




for any $x \in G$.


Taking S of size 1 allows elements that are the identity on a lower face and equal to x on the corresponding upper face for some fixed $x \in G_1$, e.g.:



Similarly, for $|S| = 2$ we get a configuration equal to x on some upper face of codimension 2, and identity elsewhere, where now $x \in G_2$ is fixed, e.g.:



Finally, we can put any element of G_3 on the topmost vertex.



Remark 2.3. At first glance, these definitions are inherently asymmetric in the sense of treating upper faces preferentially to other faces. However, it is easy to see that $[x]_F$ is in $\text{HK}^k(G_\bullet)$ for any face F of codimension r , provided $x \in G_r$.

For instance, if $S = \{i\}$ for some $i \in [k]$, and let F denote the lower face $\{\omega_i = 0\}$. Then

$$[x]_F = [x^{-1}]_{F_S} [x]_{F_\emptyset}$$

and this is clearly in $\text{HK}^k(G_\bullet)$. In general one can argue by induction on the co-dimension.

Ultimately, though, we are not interested in nilpotent groups but in nilmanifolds G/Γ , their compact homogeneous spaces. The notion of Host–Kra cubes over G goes over directly to a notion on G/Γ .

Definition 2.4. Let G_\bullet be a degree s filtered Lie group. Suppose for simplicity that G is connected.² Also, let Γ be a discrete and co-compact subgroup of G (the latter meaning the quotient G/Γ is compact). Under these hypotheses, the quotient space G/Γ (which need not be a group) is termed a *nilmanifold*.

Suppose furthermore that $\Gamma \cap G_i$ is discrete and co-compact in G_i for each $i \geq 0$. If this property holds, we say that Γ is *compatible* with the filtration. Then for each $k \geq 0$, we define the Host–Kra cubes $\text{HK}^k(G_\bullet)/\Gamma$ to be the subset of $(G/\Gamma)^{\{0,1\}^k}$ given by the image of $\text{HK}^k(G_\bullet)$ under the quotient map

$$\begin{aligned} \pi: G^{\{0,1\}^k} &\rightarrow (G/\Gamma)^{\{0,1\}^k} \\ (g_\omega)_{\omega \in \{0,1\}^k} &\mapsto (g_\omega \Gamma)_{\omega \in \{0,1\}^k} . \end{aligned}$$

Note that we abuse notation to let $\text{HK}^k(G_\bullet)/\Gamma$ denote the *pointwise* quotient by Γ rather than a conventional quotient of groups: this notation is not meant to identify Γ with a subgroup of $\text{HK}^k(G_\bullet)$. Equivalently, this quotient may be identified with $\text{HK}^k(G_\bullet) / \left(\Gamma^{\{0,1\}^k} \cap \text{HK}^k(G_\bullet) \right)$.

The topological conditions are chosen to allow the following conclusion.

Proposition 2.5. *For G_\bullet and Γ as in the above definition, the space $\text{HK}^k(G_\bullet)/\Gamma$ is a compact subset of $(G/\Gamma)^{\{0,1\}^k}$ for all $k \geq 0$.*

²In [GMV16a, GMV16b] there are good reasons to relax this assumption, but they do not apply for now.

See [GT10, Lemma E.10] for a proof.

Finally, we summarize some properties of Host–Kra cubes on nilmanifolds that will be of significance in the next section. The proofs of these properties appear in Appendix A, or follow easily from those results.

Proposition 2.6. *Let G_\bullet and $\Gamma \subseteq G$ be as in Definition 2.4, and in particular G_\bullet has degree s . Then the following hold.*

- (i) *(Symmetries) The space $\mathrm{HK}^k(G_\bullet)/\Gamma$ is invariant under a permutation of the k coordinate axes, reflecting in a coordinate axis, or any combination of these.*
- (ii) *(Compatibility) If $c \in \mathrm{HK}^k(G_\bullet)/\Gamma$ then any face of c of dimension ℓ (or more generally any “subcube”, allowing some diagonal slicing) is an element of $\mathrm{HK}^\ell(G_\bullet)/\Gamma$.*
- (iii) *(Corner constraint) Suppose $c, c' \in \mathrm{HK}^{s+1}(G_\bullet)/\Gamma$ and $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{1}$. Then $c = c'$.*
- (iv) *(Corner completion) Suppose $\lambda: \{0, 1\}^k \setminus \{\vec{1}\} \rightarrow G/\Gamma$ is a configuration such that for every lower face of $\{0, 1\}^k$ of codimension 1, the restriction of λ to that face is in $\mathrm{HK}^{k-1}(G_\bullet)/\Gamma$. Then λ can be extended to an element of $\mathrm{HK}^k(G_\bullet)/\Gamma$.*

Here, and throughout, $\vec{1}$ denotes the element $(1, 1, \dots, 1) \in \{0, 1\}^k$.

Proof. Parts (i) and (ii) follow directly from the definition. For (iii) and (iv), see Proposition A.16 and Proposition A.17. \square

3. CUBESPACES AND NILSPACES

3.1. Definitions. We now give the formal definition and axioms of *nilspaces*, and related notions. In fact, the notion of a nilspace captures several distinct hypotheses of differing strength. We will outline these individually, starting with the weakest, before amalgamating them into a final definition. We follow [HK08, ACS12] closely, although our terminology differs.

The very weakest structure we will wish to consider is termed a *cubespace*. Informally, this is just a topological space equipped with some notion of when 2^k points form a cube,³ and satisfying several fairly basic sanity conditions.

To define these sanity conditions, we need to define some nomenclature for certain maps on the discrete cube $\{0, 1\}^k$.

Definition 3.1. A map $\rho: \{0, 1\}^k \rightarrow \{0, 1\}^\ell$ is called a *morphism of discrete cubes* if it is the restriction to $\{0, 1\}^k$ of an affine-linear map $\mathbb{Z}^k \rightarrow \mathbb{Z}^\ell$.

Equivalently, this holds if and only if ρ has the form

$$\rho(x_1, \dots, x_k) = (\sigma_1(x_1, \dots, x_k), \dots, \sigma_\ell(x_1, \dots, x_k))$$

³Following [ACS12], we use the term “cube” throughout to refer to these distinguished collections of 2^k points. These objects in fact seldom resemble geometric cubes, and the term “parallelepiped” used by Host and Kra [HK08] is more accurate. However, “cube” has a significant advantage in brevity.

where each function σ_i is one of:

- identically 0;
- identically 1;
- equal to x_j for some j ;
- equal to $1 - x_j$ for some j .

It is straightforward to see that these definitions are indeed equivalent.

Informally, these morphisms of the discrete cube correspond to the fairly natural operations:

- permute the coordinates of $\{0, 1\}^k$;
- reflect in any coordinate axis;
- embed $\{0, 1\}^k$ as a “slice” in $\{0, 1\}^\ell$ for $\ell > k$;
- “project” $\{0, 1\}^k$ onto $\{0, 1\}^\ell$ for $\ell < k$ by deleting a coordinate;

and functions obtained from these by composition.


Definition 3.2. A *cubespace* is a metric space X , together with closed sets $C^k \subseteq X^{\{0,1\}^k}$ of k -cubes for each $k \geq 0$, satisfying the following condition. Suppose $\rho: \{0, 1\}^k \rightarrow \{0, 1\}^\ell$ is a morphism of discrete cubes and $c: \{0, 1\}^\ell \rightarrow X$ is in C^ℓ . Then $c \circ \rho: \{0, 1\}^k \rightarrow X$ is in C^k .

Furthermore, it is always assumed that $C^0 = X$.


We will typically abuse notation to allow X to refer either to the underlying topological space, or to the full cubespace structure (X, C^k) . We write $C^k(X)$ in place of C^k whenever there is ambiguity about which cubespace we are referring to.

Let us unpack what this means in terms of the elementary operations above.


- If $c \in X^{\{0,1\}^k}$ is a k -cube, then permuting the k coordinates gives another k -cube; e.g.:




- Similarly, reflecting in any coordinate axis gives another k -cube; e.g.:



- If $c \in X^{\{0,1\}^k}$ is a k -cube, then restricting to an ℓ -dimensional “slice” gives an ℓ -cube:



- If $c \in X^{\{0,1\}^k}$ is a k -cube, the configuration obtained by placing two copies of c adjacent to each other is a $(k+1)$ -cube, and so on:



Remark 3.3. Note that applying this last point repeatedly, we deduce that any constant configuration (i.e. $c(\omega) = x$ for all ω) is automatically a cube.

It is intuitively fairly reasonable that any well-behaved notion of “cube” analogous to parallelepipeds or the Host–Kra construction on nilmanifolds should at the very least obey these properties.

We now consider the “corner constraint” discussed previously. Recall we said informally that, for some specific dimension of cube, all but one of the vertices should determine the last one.

Definition 3.4. We say a cubespace X has k -*uniqueness* if the following holds: whenever $c, c' \in C^k(X)$ and $c(\omega) = c'(\omega)$ for all $\omega \in \{0, 1\}^k \setminus \{\vec{1}\}$ then $c = c'$.

The “dual” of this property is equally important but harder to motivate: if all but one vertex of a k -cube is specified in a consistent way, then there is *at least one* way to complete the missing vertex to give a cube.

Definition 3.5. We say a cubespace X has k -*completion* if the following holds. Suppose $\lambda: \{0, 1\}^k \setminus \{\vec{1}\} \rightarrow X$ has the property that every “lower face” is a $(k-1)$ -cube, i.e. for each $1 \leq i \leq k$ the map


$$\begin{aligned} \{0, 1\}^{k-1} &\rightarrow X \\ (\omega_1, \dots, \omega_{k-1}) &\mapsto \lambda(\omega_1, \dots, \omega_{i-1}, 0, \omega_i, \dots, \omega_{k-1}) \end{aligned}$$

is in $C^{k-1}(X)$. Then there exists some $x \in X$ such that

$$c: \{0, 1\}^k \rightarrow X$$

$$\omega \mapsto \begin{cases} \lambda(\omega) & : \omega \neq \vec{1} \\ x & : \omega = \vec{1} \end{cases}$$

is in $C^k(X)$.



We call a configuration of the form λ (with the same properties) a *k-corner*.

While uniqueness will typically be specified for one particular value of k , a well-behaved space will have k -completion for *all* $k \geq 0$.

Definition 3.6. We say a cubespace X is a *nilspace of degree s* if s is the minimal number such that X has $(s + 1)$ -uniqueness, and if X has k -completion for all k . We say it is simply a *nilspace* if it is a nilspace of degree s for some s .

Remark 3.7. One way of thinking about how the k -completion hypothesis fits into the picture is that it asserts that X is not missing any points.

Suppose X is a compact nilspace of degree s , and $S \subseteq X$ an arbitrary closed subset. We can make S into a cubespace by taking $C^k(S) = C^k(X) \cap S^{\{0,1\}^k}$, i.e. just taking those cubes of X whose vertices lie in S . We call this cubespace the subcubespace of X *induced* by S .

With this cubespace structure, it is easy to verify that $(S, C^k(S))$ is still a compact cubespace with $(s + 1)$ -uniqueness; i.e., it satisfies all the hypotheses of being a (compact) nilspace of degree s except for the k -completion one. However, $(S, C^k(S))$ will typically *not* have k -completion for any $k > 1$, unless S takes a very particular form.

So, if we were to omit the k -completion axiom from the definition of a nilspace, it would become fairly hopeless to ask for a rigid classification theorem for nilspaces, since there are a huge variety of examples of this type. By contrast, if we do impose k -completion, we will see that this rules out such examples in general and imposes significantly more rigidity on the class of nilspaces.

We will be interested not only in cubespaces and nilspaces in isolation, but also in maps between them. Informally, a map $X \rightarrow Y$ preserves the relevant structures if it takes any cube of X to a cube of Y .

Definition 3.8. Suppose X, Y are cubespaces with underlying spaces X, Y respectively. Let $\phi: X \rightarrow Y$ be a continuous map. We say ϕ is a *cubespace morphism* or just *morphism* if $\phi(C^k(X)) \subseteq C^k(Y)$ for all $k \geq 0$.

We will need one further technical definition.

Definition 3.9. A cubespace X is called *ergodic* if $C^1(X) = X^2$, i.e. if every pair (x, x') is a 1-cube. Furthermore we say it is *k-ergodic* if $C^k(X) = X^{\{0,1\}^k}$, i.e. if every configuration is a k -cube.

Note that our terminology differs slightly from that in [ACS12], in that we do not insist that a nilspace be ergodic: this is built into the definition of a nilspace in that work. As a consequence, the phrase “ergodic nilspace” will occur frequently in our statements.

It can be shown (although we will not need to do so) that a suitably nice non-ergodic space – in particular, a non-ergodic nilspace – decomposes as a disjoint union of ergodic components, whose cubespace structures essentially do not interact at all. Hence one loses almost no generality by working in the ergodic setting.

Also, we remark that [ACS12] has the separate notions of “abstract” nilspaces, which have no topology, and “compact” ones, which have topology but are always assumed to be compact. By contrast, we assume all nilspaces have topology (though it could be the discrete topology), but do not insist they be compact. This discrepancy has no particular significance, except that it allows us to consider e.g. nilpotent Lie groups as nilspaces with their natural topologies. Again, the phrase “compact, (ergodic) nilspace” will occur often in what follows.

3.2. Host–Kra nilmanifolds are nilspaces. Having motivated these nilspace axioms (and other definitions) in terms of Host–Kra cubespaces and nilmanifolds, it is reasonable to want to check that the latter are actually instances of the former.

We state some facts of this nature now. Most aspects of the proofs are straightforward; some, surprisingly, are much less so, and these are deferred to Appendix A which expounds the theory of Host–Kra cubespaces more fully than we have done so far.

Proposition 3.10. *Let G/Γ be a (filtered) nilmanifold of degree s , in the sense of the previous section, equipped with its Host–Kra cubes. Then it is a nilspace of degree s .*

If the filtration G_\bullet is proper (i.e. $G_0 = G_1$) then G/Γ is ergodic.

If G'/Γ' is another nilmanifold, and $\psi: G \rightarrow G'$ a group homomorphism such that $\psi(G_i) \subseteq G'_i$ for all i and $\psi(\Gamma) \subseteq \Gamma'$, then the induced map $\phi: G/\Gamma \rightarrow G'/\Gamma'$ is a cubespace morphism.

Similarly, for any $x \in G_0$, $g \in G_1$, the map $n \mapsto g^n x \Gamma$ is a cubespace morphism $\mathbb{Z} \rightarrow G/\Gamma$.

Here, as throughout these papers, \mathbb{Z} is considered as a filtered group with filtration $\mathbb{Z}_0 = \mathbb{Z}_1 \supseteq \{0\}$, and as a cubespace by setting $C^k(\mathbb{Z}) = \text{HK}^k(\mathbb{Z}_\bullet)$.

Proof. The fact that G/Γ is a nilspace is essentially a restatement of Proposition 2.6.

The ergodicity statement is immediate from the definition of $\text{HK}^1(G_\bullet)$.

The fact that a group homomorphism $\psi: G \rightarrow G'$ taking G_i into G'_i maps $\text{HK}^k(G_\bullet)$ into $\text{HK}^k(G'_\bullet)$ is straightforward by considering the images of the generators. Recovering the corresponding statement for the quotients G/Γ , G'/Γ' is routine.

Finally, we note that for any $x \in G_0$, $g \in G_1$, $a \in \mathbb{Z}$ and $h \in \mathbb{Z}^k$, the configuration

$$\begin{aligned} \{0, 1\}^k &\rightarrow G \\ \omega &\mapsto g^{a+\omega \cdot h} x \end{aligned}$$

is in $\text{HK}^k(G_\bullet)$, since it is a product of $(k+1)$ generators:

$$[g^{h_1}]_{F_{\{1\}}} \cdots [g^{h_k}]_{F_{\{k\}}} \cdot [g^a x]_{F_\emptyset},$$

and the last part follows. \square

We have been light on detail in this proof: partly because this result is not logically necessary in what follows. However, working through the details is an excellent exercise for the reader unfamiliar with either these definitions or the machinery of Host–Kra cubes (or both).

3.3. High-dimensional cubes. One slightly unsatisfactory feature of our definitions is that the data of a cubespace involves spaces $C^k(X)$ for infinitely many $k \geq 0$. We are used to considering cubes or parallelepipeds of only bounded dimension in any given problem (e.g. of dimension 3 when considering the U^3 norm).

In fact, under good hypotheses we can see that for sufficiently large k the data of $C^k(X)$ contains no new information.

As one further reassuring sanity-check, and to see these definitions in action, we state this now.

Proposition 3.11. *Let $s \geq 1$ be an integer, X a nilspace of degree s , and $k \geq s+1$. Then a configuration $c: \{0, 1\}^k \rightarrow X$ is in $C^k(X)$, if and only if every face of c of dimension $(s+1)$ is in $C^{s+1}(X)$.*

Proof. The “only if” direction is direct from the cubespace axioms. For the “if” direction, we argue by induction on k , the base case $k = s+1$ being trivial.

Given such a $c: \{0, 1\}^k \rightarrow X$, consider the restriction $c|_{\{0, 1\}^k \setminus \{\vec{1}\}}$. By inductive hypothesis it is a k -corner. Hence, there is a cube $c' \in C^k(X)$ such that $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{1}$. But now, considering any upper face c' of dimension $(s+1)$ and the corresponding one of c , and invoking $(s+1)$ -uniqueness, we see that $c(\vec{1}) = c'(\vec{1})$ and hence c is a k -cube. \square

4. THE STRUCTURE THEOREM

4.1. Statement. We are now in a position to state the main structural result. This result is essentially due to Antolín Camarena and Szegedy [ACS12, Theorems 4 and 7]; however, our formulation is somewhat stronger. We point out the differences below.

Theorem 4.1. *Let $X = (X, C^k)$ be a compact ergodic nilspace of degree s . Further suppose that the spaces $C^k(X)$ are connected for each $k \geq 0$; so in particular $X = C^0(X)$ is connected.*

Then X is an “inverse limit of nilmanifolds” in the following sense. There exists:

- *a sequence $G^{(n)}$ of connected Lie groups equipped with filtrations $G_{\bullet}^{(n)}$ of degree at most s (with $G_i^{(n)}$ also connected for each i);*
- *discrete co-compact subgroups $\Gamma^{(n)}$ of $G^{(n)}$ such that $\Gamma^{(n)} \cap G_i^{(n)}$ is discrete and co-compact in $G_i^{(n)}$ for each i ; and*
- *surjective group homomorphisms $\phi_{n,m}: G^{(n)} \rightarrow G^{(m)}$ for each $n \geq m$, such that $\phi_{n,m}(G_i^{(n)}) = G_i^{(m)}$ for each $i \geq 0$ and also $\phi_{n,m}(\Gamma^{(n)}) \subseteq \Gamma^{(m)}$, so that we get a map of nilmanifolds $G^{(n)}/\Gamma^{(n)} \rightarrow G^{(m)}/\Gamma^{(m)}$ that sends cubes (surjectively) to cubes;*

such that

- *$X = \varprojlim (G^{(n)}/\Gamma^{(n)})$ as a topological space, and*
- *the cubes $C^k(X)$ coincide with the inverse limit of Host–Kra cubes; i.e., writing $\pi_n: X \rightarrow (G^{(n)}/\Gamma^{(n)})$ for the projection map arising from the above inverse limit, we have*

$$C^k(X) = \bigcap_n \pi_n^{-1} \left(C^k \left(G^{(n)}/\Gamma^{(n)} \right) \right) .$$

Remark 4.2. The last two conditions, together with the fact that $\phi_{n,m}$ are cubespace morphisms between nilmanifolds (see Proposition 3.10), say in some sense that X is an “inverse limit of nilmanifolds, in the category of cubespaces”. This is roughly the formulation one obtains from [ACS12], combining [ACS12, Theorems 4 and 7]. (In fact the results from [ACS12] give slightly more detail than that, specifically that the maps $\phi_{n,m}$ are “fiber-surjective”, or in our terminology “fibrations”; see Definition 7.1 and Remark 7.4. This same condition comes out of [GMV16b, Theorem 1.26], which is equivalent to [ACS12, Theorem 4].)

However, the conclusion of Theorem 4.1 is stronger than this, in that it gives even more information about the maps $\phi_{n,m}$: namely, that they come from group homomorphisms $G^{(n)} \rightarrow G^{(m)}$, and moreover that these group homomorphisms are surjective on the filtrations in the obvious sense. Neither of these properties holds in general for a cubespace morphism between nilmanifolds.

Remark 4.3. We call a cubespace satisfying the hypothesis that $C^k(X)$ is connected for all k *strongly connected*. Without some similar connectivity hypothesis, no complete structure theorem is currently available in general.

Using Theorem 5.12 (see below) or equivalently [ACS12, Theorem 4], one can get some partial information by identifying a compact nilspace with an inverse limit of “finite rank” nilspaces, which

are in particular e.g. topological manifolds of finite dimension. However, the structure of such spaces remains unclear: they need not necessarily be nilmanifolds. Some structural results in the special case of finite nilspaces appear in [Sze10].

The analogous result implied by [ACS12] has this strong connectivity hypothesis replaced by an assumption on the “structure groups” of X (see below). That assumption easily implies strong connectivity (see [GMV16a, Proposition 2.4]), but for the converse, we need the full force of our structure theorem.

For clarity, we summarize the dictionary between statements in these papers and those in [ACS12] in Section 5.4 below, once we have introduced the relevant concepts.

Remark 4.4. It is reasonably straightforward to argue that any cubespace X that is an inverse limit of nilmanifolds, subject to some further technical hypotheses on this inverse limit, satisfies the hypotheses. Hence in some sense this result is best possible.

One can certainly not remove the inverse limit, as can be seen by considering any compact connected abelian group that is not a Lie group: such an object is a nilspace of degree 1 but is not a degree 1 nilmanifold, as the latter are all just tori.

It is also possible to find examples of nilspaces satisfying the hypotheses for $s = 2$, that are not an homogeneous space of any group. Such an example is due to Rudolph [Rud95]. We will not pursue this here, although related ideas may be the subject of a future paper by the authors.


4.2. The case $s = 1$. To provide some motivation for this structural result, we briefly comment on the case $s = 1$. In this setting, Theorem 4.1 is much easier.

Proposition 4.5. *A compact ergodic nilspace X of degree 1 naturally has the structure of a compact abelian group.*

Proof sketch. We give only a partial account of the proof, since these ideas are covered in detail in Section 6.4.

Essentially our task is to use the data we have – namely, unique corner-completion of 2-cubes – to recover a group operation on X .

There is only really one thing to try here, and it works. Fix any element $e \in X$, which will be our identity element. Given $x, y \in X$ consider the 2-corner:



and define $x + y$ to be the element obtained by completing it, i.e. $x + y = *$. We can define inverses by a similar process, and define the identity to be e .

It now suffices to check that this defines a topological abelian group operation. Many of the corresponding properties are immediate from the nilspace axioms. The hardest part is checking associativity, which requires the completion property on 3-dimensional cubes; we omit this part. \square

The remaining content of Theorem 4.1 in this case is that any connected compact abelian group is an inverse limit of tori, which is a classical result.

4.3. Steps in the proof. The proof of Theorem 4.1 splits fairly cleanly into three parts.

The first is a weaker, fairly elementary structure theorem. Roughly speaking, this says that a nilspace of degree s has the structure of a tower of extensions $X \rightarrow X_{s-1} \rightarrow \cdots \rightarrow X_0 = \{*\}$ where each of the fibers is a compact abelian group. We refer to this as the *weak structure theorem*.

The second says that, under some strong technical assumptions, a compact ergodic strongly connected nilspace of degree s actually *is* a nilmanifold G/Γ with usual Host–Kra cubes (and a degree s filtration). These technical assumptions are phrased in terms of the compact abelian groups appearing in the weak structure theorem; specifically, that they are (abelian) Lie groups, i.e. $(\mathbb{R}/\mathbb{Z})^d \times K$ for some $d \geq 0$ and K finite. This is the step where a non-abelian group operation is recovered from just the cubespace structure, the key step being to consider the automorphism group of the cubespace X , and to show that it is large enough in a certain sense.

The final step requires us to show that a cubespace satisfying the hypotheses of Theorem 4.1 can be written as an inverse limit of ones which satisfy the strong technical assumptions alluded to in part two. Given the previous discussion, this is very closely related to the statement that a connected compact abelian group is an inverse limit of tori, and this fact is a key ingredient in the proof. However, bolting everything together is one of the more technically difficult aspects of the whole argument.

Only the first part of this plan will be proven fully in this paper, although we will give rather fuller outlines of the latter two below. These are proven fully in two further papers by the authors, the second part in [GMV16a] and the third in [GMV16b].

5. A MORE DETAILED SUMMARY OF THE ARGUMENT

We will now expand on the outline given above of each of the three stages in the proof of Theorem 4.1.

5.1. An outline of the weak structure theory. The simplest examples of nilspaces are compact abelian groups A together with their Host–Kra cubes (with the degree 1 filtration $A = A_0 = A_1 \supseteq A_2 = \{0\}$). The k -cubes of these spaces are precisely the k -dimensional parallelepipeds on A in the usual sense:

$$\{0, 1\}^k \rightarrow A$$

$$\omega \mapsto x_0 + \sum_{i=1}^k \omega_i x_i$$

where $x_0, \dots, x_k \in A$.

The simplest higher-degree examples come from the Host–Kra construction on A (still a compact abelian group) given the degree s filtration

$$A = A_0 = A_1 = \dots = A_s \supseteq A_{s+1} = \{0\}$$

where cubes are given by $C^k(A) = \text{HK}^k(A_\bullet)$. We denote this cubespace by $\mathcal{D}_s(A)$. It is easy to check (or to deduce from general considerations) that $\mathcal{D}_s(A)$ is a compact, ergodic (and in fact s -ergodic) nilspace of degree s .

In fact there are several equivalent ways to define $\mathcal{D}_s(A)$. We briefly note some equivalent characterizations.

Proposition 5.1. *If $c: \{0, 1\}^{s+1} \rightarrow A$ is a configuration, then $c \in C^{s+1}(\mathcal{D}_s(A))$ if and only if*

$$\sum_{\omega \in \{0, 1\}^{s+1}} (-1)^{|\omega|} c(\omega) = 0.$$

Moreover, for any $k \geq 0$ and $c: \{0, 1\}^k \rightarrow A$, we have $c \in C^k(\mathcal{D}_s(A))$ if and only if every face of c of dimension $(s+1)$ is a cube; or, if and only if $c \circ \eta \in C^{s+1}(\mathcal{D}_s(A))$ for every morphism of discrete cubes $\eta: \{0, 1\}^{s+1} \rightarrow \{0, 1\}^k$.

Proof. The first claim is proved in [A.9](#).

It is clear from the cubespace axioms that if $c \in C^k(\mathcal{D}_s(A))$ and $\eta: \{0, 1\}^{s+1} \rightarrow \{0, 1\}^k$ is a morphism of discrete cubes then $c \circ \eta \in C^{s+1}(\mathcal{D}_s(A))$, and this includes the case of faces of c .

Conversely, by [Proposition 2.6](#), we know that $\mathcal{D}_s(A)$ is a nilspace of degree s . So if $c: \{0, 1\}^k \rightarrow A$ and every face of c of dimension $(s+1)$ is a cube, then c is a cube by [Proposition 3.11](#). \square

The weak structure theorem states that, while these are certainly not the only examples of nilspaces, any (compact, ergodic) nilspace X of degree s can be *built up* from spaces of the form $\mathcal{D}_k(A_k)$ for some compact abelian groups A_k for $1 \leq k \leq s$. Indeed, X can be realized as a tower of extensions

$$X \rightarrow \pi_{s-1}(X) \rightarrow \pi_{s-2}(X) \rightarrow \dots \rightarrow \pi_0(X) = \{*\}.$$

such that $\pi_i(X)$ is a nilspace of degree i and the fibers of the map $\pi_k(X) \rightarrow \pi_{k-1}(X)$ admit a free transitive action of a compact abelian group A_k through which the cubespace structure on the fiber can be identified by $\mathcal{D}_k(A_k)$.

Before we can formulate these results more precisely we need to discuss how we can put a cubespace structure on a quotient of a cubespace.

Definition 5.2. Let $(X, C^k(X))$ be a compact cubespace, and suppose \sim is a closed equivalence relation on X . The quotient space X/\sim (with the quotient topology) is considered to be a cubespace, by setting $C^k(X/\sim)$ to be $C^k(X)/\sim$, i.e. the image of $C^k(X)$ under pointwise application of \sim . In other words this means that a configuration in X/\sim is a cube if and only if there is some representative of it that is a cube in X .

It is straightforward from the definition that this indeed defines a cubespace structure. In the next definition, we specify the equivalence relation that give rise to the quotients $\pi_i(X)$. This may be compared with [ACS12, Definition 2.3].

Definition 5.3. For any $k \geq 0$, we define \sim_k on a (compact, ergodic) cubespace X with n -completion for all n , by $x \sim_k y$ if and only if there exist $c, c' \in C^{k+1}(X)$ such that $c(\vec{1}) = x$, $c'(\vec{1}) = y$ and $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{1}$. We call \sim_k the k -th canonical equivalence relation. We denote by π_k the quotient map $X \rightarrow X/\sim_k$ and call it the k -th canonical projection.

If \sim is an equivalence relation such that the factor X/\sim has $(s+1)$ -uniqueness, then it must contain \sim_s , which is the motivation for the definition. However, it requires proof that \sim_s is indeed an equivalence relation and the quotient is indeed a nilspace. We will discuss these facts and further properties of \sim_s in Section 6.3.

We are now ready to state the weak structure theorem, which is proved in Sections 6 and 7. See also [ACS12, Theorem 1], and also [HK08, Section 5] for a related discussion.

Theorem 5.4 (Weak Structure Theorem). *Let X be a compact ergodic nilspace of degree s for some $s \geq 1$. Then there exists a compact abelian group $A = A_s(X)$, the “ s -th structure group” of X , acting continuously on X , such that:*

- (i) *the action of A is free on X , and its orbits are precisely the fibers of $X \rightarrow \pi_{s-1}(X)$;*
- (ii) *this induces a pointwise action of $C^k(\mathcal{D}_s(A))$ on $C^k(X)$: again this action is free, and its orbits are precisely the fibers of the induced map $C^k(X) \rightarrow C^k(\pi_{s-1}(X))$.*

We define the k -th structure group $A_k(X)$ of X as the group $A_k(\pi_k(X))$ that arises when we apply this theorem to the nilspace $\pi_k(X)$ (with $s = k$). We will see in the proof that these groups are defined canonically in terms of the cubespace structure.

We note that the weak structure theorem allows the following partial reconstruction of cubes on X . Let $c \in C^k(\pi_{s-1}(X))$ be a cube. If we know that $\tilde{c} \in C^k(X)$ is a cube that projects to c (that is $\pi_{s-1}(\tilde{c}) = c$), then we can find all cubes projecting to \tilde{c} . Indeed, let $c' : \{0, 1\}^k \rightarrow X$ be a configuration with $\pi_{s-1}(c') = c$, and write $a : \{0, 1\}^k \rightarrow A_s(X)$ for the unique configuration such that $a(\omega) \cdot \tilde{c}(\omega) = c'(\omega)$. Then c' is a cube if and only if $a \in C^k(\mathcal{D}_s(A_s))$. However, the weak structure theorem provides no insight for finding the first lift \tilde{c} of a cube in $C^k(\pi_{s-1}(X))$.

Remark 5.5. We note that no topological input is really used in this theorem, except to show topological conclusions. For instance, if X is a possibly infinite *discrete* nilspace (i.e. ignoring topology), the same proofs will apply, resulting in a discrete abelian group A acting on the space and so on.

We will not consider the result in such generality, as we have no need for such statements in applications.

5.2. Automorphisms and recovering a nilmanifold structure. Stage II in our plan for proving Theorem 4.1 was to establish a strong structural result, stating that X actually is a nilmanifold G/Γ ,

under some additional technical assumptions. We state this result, which is a variant of [ACS12, Theorem 7], now.

Theorem 5.6. *Suppose X is a compact, ergodic, strongly connected nilspace of degree s . Further suppose that the structure groups A_t for $1 \leq t \leq s$ are all (compact abelian) Lie groups, i.e. isomorphic to tori $(\mathbb{R}/\mathbb{Z})^d \times K$ for some $d = d(A_t) \geq 0$ and some finite K .*

Then X is isomorphic to a nilmanifold G/Γ , equipped with its Host–Kra cubes. That is, there exists a connected Lie group G equipped with a filtration of degree s (and such that G_i is also connected for each i), and a discrete and co-compact subgroup Γ of G compatible with the filtration (cf. Definition 2.4), such that

- X is homeomorphic to G/Γ , and
- under this homeomorphism, $C^k(X)$ is identified with $\mathrm{HK}^k(G_\bullet)/\Gamma$.

A complete proof of this theorem can be found in another paper of the authors [GMV16a, Theorem 2.18].

Remark 5.7. The difference between Theorem 5.6 and [ACS12, Theorem 7] is that the latter makes the assumption that the structure groups are tori $\mathbb{R}^d/\mathbb{Z}^d$ – i.e. that the finite group K is necessarily trivial, or equivalently that $A_k(X)$ is connected – in place of strong connectivity.

As we mentioned in Remark 4.3, these two conditions are equivalent, though one of the implications is a non-trivial fact. It will be slightly convenient for the current exposition to assume the hypothesis about connectivity of the structure groups from the start. This simplifies, but by no means trivializes, the proof.

Remark 5.8. It is possible to formulate a version of Theorem 5.6 without any reference to the structure groups. Then one needs to add some hypotheses on the topology of the space X , e.g. that it is locally connected and has finite Lebesgue covering dimension. However, the proof of this statement requires the general structure theorem discussed in Section 4. For details, we refer to the Appendix of the companion paper [GMV16b].

At the heart of the challenge in proving the above theorem is recovering the non-abelian group operation on G . Although we were able to find a very explicit and combinatorial construction of the group operations on the structure groups A_t just by exploiting concatenation and cube completion, there are seemingly fatal obstacles to such a simple approach working for the non-abelian group G . Essentially this is because Theorem 5.6 is simply not true, even in spirit, if the strong connectivity assumptions on X are dropped, and hence any proof must have some non-elementary topological aspect.

Instead, the group operation is recovered in a completely different way: by considering the group of automorphisms of a nilspace X .

Definition 5.9. Write $\mathrm{Aut}(X)$ for the (topological) group of cubespace automorphisms of X ; that is, of homeomorphisms $\phi: X \rightarrow X$ such that both ϕ and ϕ^{-1} are cubespace morphisms.

The topology being used here is compact-open. It turns out that $\text{Aut}(X)$ has attached to it a canonical filtration,

$$\text{Aut}(X) = \text{Aut}_0(X) \supseteq \text{Aut}_1(X) \supseteq \cdots \supseteq \text{Aut}_s(X) \supseteq \{\text{id}\}$$

of degree s , assuming X is a nilspace of degree s . The definition of the filtration is very explicit and depends only on the cubes of X , but we will not go into the details here.

It will typically be the case that the first containment $\text{Aut}_0(X) \supseteq \text{Aut}_1(X)$ is strict, and so we have no guarantee that $\text{Aut}(X)$ is nilpotent – in fact, often it will not be.⁴ For this reason it is better to work with the subgroup $\text{Aut}_1(X)$, sometimes called the group of 1-*translations*. This inherits a filtration

$$\text{Aut}_1(X) = \text{Aut}_1(X) \supseteq \cdots \supseteq \text{Aut}_s(X) \supseteq \text{Aut}_{s+1}(X) = \{\text{id}\}$$

which guarantees that $\text{Aut}_1(X)$ is nilpotent of nilpotency class at most s .

In general, the group $\text{Aut}_1(X)$ need not be connected. It turns out to be vital to work instead with the identity component of $\text{Aut}_1(X)$, which we term $\text{Aut}_1^\circ(X)$, and similarly with the revised filtration $\text{Aut}_i^\circ(X)$, the identity component of $\text{Aut}_i(X)$. Note this need *not* be the same as $\text{Aut}_1^\circ(X) \cap \text{Aut}_i(X)$, and hence one should check that this is actually a filtration, but this is not hard to show.

Given this filtration, we can consider the Host–Kra cube-groups $\text{HK}^k(\text{Aut}_1(X)_\bullet)$ or $\text{HK}^k(\text{Aut}_1^\circ(X)_\bullet)$. Again these have a natural interpretation in terms of the cubes of X . One consequence is that if $x_0 \in X$ and $(\phi_\omega)_{\omega \in \{0,1\}^k} \in \text{HK}^k(\text{Aut}_1(X)_\bullet)$ then the configuration $(\phi_\omega(x_0))_{\omega \in \{0,1\}^k}$ is a cube in $C^k(X)$.

In other words given $x_0 \in X$, the map

$$\begin{aligned} \text{Aut}_1(X) &\rightarrow X \\ \phi &\mapsto \phi(x_0) \end{aligned}$$

is a cubespace morphism if $\text{Aut}_1(X)$ is given its Host–Kra cubes (and hence this restricts to a cubespace morphism on $\text{Aut}_1^\circ(X)$). It will not be an isomorphism, as it is typically not injective: x_0 may have non-trivial stabilizer. But we can define another map of cubespaces

$$\begin{aligned} \text{Aut}_1^\circ(X) / \text{Stab}(x_0) &\rightarrow X \\ \phi \text{ Stab}(x_0) &\mapsto \phi(x_0) \end{aligned}$$

which is well-defined and again a cubespace morphism.

The hope is that this is in fact surjective, and furthermore an isomorphism of cubespaces (which is a strictly stronger property). If in addition $\text{Aut}_1^\circ(X)$ were a Lie group and $\text{Stab}(x_0)$ a discrete and co-compact subgroup, we would have identified X with a nilmanifold, as required.

However, we have not yet even established that $\text{Aut}(X)$ is not just the trivial group, in which case this map would be very far from surjective. Establishing that this is not the case, and that X has

⁴This can be seen even when $X = (\mathbb{R}/\mathbb{Z})^2$ with the degree 1 filtration: $\text{Aut}(X)$ contains a copy of $\text{SL}_2(\mathbb{Z})$.

“enough” automorphisms for this strategy to succeed, requires a reasonable proportion of the hard work of this whole approach.

We will offer a few further clues as to the structure of this part of the proof. Suppose for now that $s = 2$, and so X is a nilspace of degree 2. Note that – contrary to what we have just pessimistically stated – we do have *some* examples of non-trivial automorphisms of X : namely, the action of the top structure group $A_2(X)$ described above. Of course, since this action fixes fibers of the quotient $X \rightarrow \pi_1(X)$ this falls a long way short of showing surjectivity.

However, we similarly have plenty of automorphisms of $\pi_1(X)$, given by the action of the next structure group $A_1(X)$. In fact (assuming X is ergodic) this action is simply transitive.

So, we know that we can move by elements of $\text{Aut}_1^\circ(X)$ along fibers of π_1 ,⁵ and also that we can move between fibers by automorphisms of $\pi_1(X)$. What we want, though is that these elements of $\text{Aut}_1^\circ(\pi_1(X))$ can be lifted to genuine automorphisms of X itself.

It is not hard to show the following, although currently we do not have the definitions in place to do so.

Lemma 5.10. *Elements of $\text{Aut}_1(X)$ commute with π_{s-1} , and so there is a natural group homomorphism $\text{Aut}_1(X) \rightarrow \text{Aut}_1(\pi(X))$. Hence similarly we get a group homomorphism $\text{Aut}_1^\circ(X) \rightarrow \text{Aut}_1^\circ(\pi(X))$.*

The whole problem therefore reduces to proving the following.

Proposition 5.11. *Suppose X is an ergodic nilspace of degree s , whose structure groups A_t are tori. Then the map $\text{Aut}_1^\circ(X) \rightarrow \text{Aut}_1^\circ(\pi(X))$ from the previous proposition is surjective.*

In the $s = 2$ case, this would tell us that $\text{Aut}_1^\circ(X)$ acts transitively on the fibers of π_1 (i.e. given any two fibers, we can map some point on one fiber to some point on the other); and since we already know that A_2 acts transitively on each fiber, this would show that $\text{Aut}_1^\circ(X)$ acts transitively on X and so the map $\text{Aut}_1^\circ(X) \rightarrow X$ is surjective.

Moving from surjectivity to isomorphism (and from the $s = 2$ case to the general case) involves some extra work, but it is mainly technical.

To prove the proposition, one begins by showing that if a 1-translation on $\pi_1(X)$ is a small perturbation of the identity, then it lifts to a 1-translation on X (if X is a nilspace of degree 2). The proof of this is roughly to relate failure of lifting to some kind of “cocycle”, and therefore express obstructions to lifting in terms of some kind of “cohomology” (where we use the term in a very loose sense). One then argues that this “cohomology” is somehow discrete, and so provided the 1-translation is small enough, no obstructions arise.

These notions of “cocycles” and “cohomology” are useful tools, and we will allude to them again below.

⁵Note we have used here the extra assumption that $A_2(X)$ is connected.

The proof of Theorem 5.6 is very closely modeled on the arguments of Antolín Camarena and Szegedy [ACS12]. The main difference is the following. In constructing the lift of a translation on the canonical quotient, we first construct (in the paper [GMV16a]) a continuous lift, which is not necessarily a translation. Then we use the action of the structure group to “correct it” to a translation. On the other hand in [ACS12], a measurable lift is constructed first, which is shown to be continuous a posteriori.

5.3. The inverse limit statement. Given this structure theorem for (ergodic, compact, strongly connected) nilspaces whose structure groups are Lie, our remaining task is to deduce something for general (ergodic, compact, strongly connected) nilspaces.

Essentially this task reduces to showing some version of the following (see also [ACS12, Theorem 4]).

Theorem 5.12. *Let X be an ergodic compact nilspace of degree s .*

Then there exists a sequence X_n of ergodic compact nilspaces of degree s , such that $A_t(X_n)$ are Lie groups for all n and all $1 \leq t \leq s$, and such that $X \cong \varprojlim X_n$ (with the inverse limit being in the sense of cubespaces).

In fact the result as proven in [GMV16b, Theorem 1.26] is marginally stronger, but in ways that use terminology we have not yet discussed.

Note that this gives some partial information about general (ergodic, compact) nilspaces without any connectivity assumption.

This is not quite enough to deduce Theorem 4.1 as we do not get the precise description of the maps appearing in the inverse limit. A complete proof of a suitably stronger result can be found in another paper of the authors [GMV16b, Theorems 1.27 and 1.28]. We give an outline below, ignoring these finer points.

We will assume again for simplicity that the assumption that X is strongly connected has been replaced by the stronger one that the structure groups $A_k(X)$ are connected for all $k \geq 1$.

The starting point is again the weak structure theorem, together with the result that any compact connected abelian group is an inverse limit of tori.

Again consider first the case $s = 2$. Since $A_2(X)$ is a compact connected abelian group, we may write $A_2(X) = \varprojlim K_m$ where K_m are tori (with given surjective maps between them). Furthermore, it is straightforward to quotient X by the action of a subgroup of A_2 (the kernel of one of the projections $A \rightarrow K_m$) to deduce the following:

Fact. *Under these hypotheses, X is an inverse limit of nilspaces X_m of degree 2, with the property that $A_2(X_m)$ is a torus for all m .*

This will allow us to assume – given a bit of work – that $A_2(X)$ is itself a torus, as otherwise we can reduce to that case.

Meanwhile, $A_1(X)$ is also a connected compact abelian group so we can write $A_1(X) = \varprojlim M_r$ for some tori M_r . We can deduce that $\pi_1(X)$ is an inverse limit of tori (in the sense of cubespaces).

The remaining challenge is to prove something like the following.

Lemma 5.13. *Let X be as above. Then for all sufficiently large integers r , there exists a canonical (degree 2, ergodic, compact) nilspace X_r and a quotient map $X \rightarrow X_r$ such that $A_2(X_r) \cong A_2(X)$, $A_1(X_r) \cong M_r$ and the diagram*

$$\begin{array}{ccc} X & \longrightarrow & X_r \\ \downarrow \pi_1 & & \downarrow \pi_1 \\ \mathcal{D}_1(A_1(X)) & \xrightarrow{\phi} & \mathcal{D}_1(M_r) \end{array}$$

commutes, where $\phi: A_1(X) \rightarrow M_r$ is the map from the inverse limit.

The problem is essentially one of “pushing forward” the degree 2 cubespace structure on X along the map ϕ . Unfortunately this operation does not make sense in general.

Again, though, it will turn out that the obstructions to building such a push-forward space are measured by a kind of cohomology; in fact the same one as was used to prove the toral structure theorem. Again, we use a kind of discreteness result for this cohomology to argue that, if M_r is sufficiently close to $\pi_1(X)$ in some sense, then these obstructions do not arise.

One way of measuring this “closeness” is the statement that the fibers of the map $\pi_1(X) \rightarrow M_r$ should have small diameter with respect to the metric on $\pi_1(X)$. Since $\pi_1(X) = \varprojlim M_r$ holds in the sense of metric spaces, this will be true provided r is sufficiently large.

Again, the general case follows a very similar pattern, with this same argument repeated once for each $1 \leq t \leq s$.

As before our approach has much in common with that in [ACS12], but differs in some important respects. Notably, the argument in that paper proceeds by establishing a correspondence between extensions a given nilspace X by a given compact abelian group A , up to isomorphism, and classes of measurable cocycles on X . An intricate argument is required to recover a topological object (the extended nilspace) from the measurable data of the cocycle. Some relative cocycle theory is then used to “push forward” the cocycle arising from $X \rightarrow \pi_1(X)$ onto $\mathcal{D}_1(M_r)$ to build X_r . By contrast, our argument realizes X_r as an explicit quotient of X , although many related tools are needed in the process.

5.4. A dictionary of statements. For clarity and convenience, we summarize which of the main statements in these papers correspond to which in [ACS12].

- The result concerning nilspaces whose structure groups are Lie, Theorem 5.6, is analogous to [ACS12, Theorem 7]; see Remark 5.7 for a detailed comparison.

- The inverse limit theorem, Theorem 5.12, is stated identically to [ACS12, Theorem 4] (though the proof strategies differ, as remarked above).
- Our main structure theorem, Theorem 4.1, has no explicit counterpart in [ACS12], but should be compared to what one obtains by concatenating [ACS12, Theorems 4 and 7]; see Remarks 4.2 and 4.3.

Theorem 4.1 also depends on [GMV16b, Theorem 1.27], to obtain the sharper inverse limit statement. The latter has no counterpart in [ACS12].

6. THE WEAK STRUCTURE THEORY

We now turn to detailed statements and proofs of the weak structure theory.

In fact, we will approach this in two stages. This section will present the “standard” weak structure theory that appears in [ACS12], and give most of the proofs (although these may differ in some places from those in [ACS12]). However, elsewhere in the project we will need a “relative” analogue of all of these statements. This relative theory is developed in Section 7; many of the proofs there will be closely analogous to those from this section.

The relative versions are a strict generalization of the non-relative ones, so logically the results in Section 7 suffice. Sometimes, therefore, we omit part of a non-relative proof and refer to the corresponding result from Section 7. Usually we do not, but this is for essentially pedagogical reasons.

6.1. Glueing. As a preliminary, we introduce one further property of cubespaces.

Definition 6.1. We say a cubespace X has the *glueing property* if “glueing” two cubes along a common face yields another cube.

Formally, suppose $c, c' \in C^k(X)$, and $c(\omega 1) = c'(\omega 0)$ for all $\omega \in \{0, 1\}^{k-1}$. (Here we use $\omega 0$ to denote $(\omega_1, \dots, \omega_{k-1}, 0)$ and so on.)

Then the configuration

$$c'' : \omega \mapsto \begin{cases} c(\omega) & : \omega_k = 0 \\ c'(\omega) & : \omega_k = 1 \end{cases}$$


is in $C^k(X)$.

The reason we have not defined this previously is that it follows from k -completion (for all k).

Proposition 6.2. *Suppose a cubespace $(X, C^k(X))$ has k -completion for all k . Then it satisfies the glueing axiom.*

Proof sketch. We will defer a complete proof until Section 7, by which time we will have developed the machinery for a clean argument. Meanwhile, we will draw some pictures in the case $k = 2$.

The idea is to place the two 2-cubes to be glued as faces of a partial 3-cube configuration; complete that (in two steps) to a 3-cube; and then observe that the “glued” configuration is a sub-cube of this, and hence a cube by the axioms.



□

6.2. The Heisenberg example. As motivation, we will examine an example called the Heisenberg nilmanifold. This is the nilspace $\text{HK}(\mathcal{H}_\bullet)/\Gamma$ constructed as in Appendix A, where

$$\mathcal{H} = \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in \mathbb{R} \right\}$$

is the Heisenberg group equipped with the filtration $\mathcal{H} = \mathcal{H}_0 = \mathcal{H}_1 \supseteq \mathcal{H}_2 \supseteq \{\text{id}\}$ where \mathcal{H}_2 is the center,

$$\mathcal{H}_2 = \left\{ \begin{pmatrix} 1 & 0 & z \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} : z \in \mathbb{R} \right\}$$

and Γ is the discrete co-compact subgroup that consists of the elements of \mathcal{H} with integral entries.

Note that, for any filtered group, G_j is a normal subgroup of G (since $[G_0, G_j] \subseteq G_j$ by the filtration property), and so we may consider the group quotient $G \rightarrow G/G_j$. This comes with an induced filtration $(G_i/G_j)_{i \geq 0}$ of degree $j - 1$.

For \mathcal{H} , the non-trivial case is $\mathcal{H}/\mathcal{H}_2$. This is isomorphic to the abelianization \mathbb{R}^2 with the degree 1 filtration $\mathbb{R}^2 = \mathbb{R}^2 \supseteq \{0\}$.

We also get an induced map on the nilmanifold \mathcal{H}/Γ , i.e. $\pi: \mathcal{H}/\Gamma \rightarrow (\mathcal{H}/\mathcal{H}_2)/(\Gamma/(\Gamma \cap \mathcal{H}_2)) = \mathbb{R}^2/\mathbb{Z}^2$. This map corresponds to

$$\pi: \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \Gamma \mapsto (x, y) \bmod 1 \in \mathbb{R}^2/\mathbb{Z}^2$$

and one can check this is well-defined. Hence, the nilmanifold \mathcal{H}/Γ has a quotient isomorphic to a torus $\mathbb{R}^2/\mathbb{Z}^2$, sometimes called the “horizontal torus”. Moreover, it is essentially automatic that this map is well-behaved with respect to the Host–Kra cubes $\text{HK}^k(\mathcal{H}_\bullet)/\Gamma$: the image of these cubes is precisely $\text{HK}^k(\mathbb{R}^2/\mathbb{Z}^2)$.

The fibers of this quotient map are of the form

$$\left\{ \begin{pmatrix} 1 & x & z' \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \Gamma : z' \in \mathbb{R}/\mathbb{Z} \right\}$$

i.e. each fiber has an action by the center \mathcal{H}_2 , or more precisely a simply transitive action by $\mathcal{H}_2/(\mathcal{H}_2 \cap \Gamma) = \mathbb{R}/\mathbb{Z}$. Again, there is a pleasant statement about cubes: e.g. given an element

$$c(\omega) = \begin{pmatrix} 1 & x_\omega & z_\omega \\ 0 & 1 & y_\omega \\ 0 & 0 & 1 \end{pmatrix} \Gamma$$

in $\mathrm{HK}^3(\mathcal{H}_\bullet)/\Gamma$, and another configuration

$$c'(\omega) = \begin{pmatrix} 1 & x_\omega & z'_\omega \\ 0 & 1 & y_\omega \\ 0 & 0 & 1 \end{pmatrix} \Gamma$$

so $\pi(c) = \pi(c')$, then one can check that c' is a cube if and only if

$$z_{000} - z_{001} - \cdots + z_{110} - z_{111} = z'_{000} - z'_{001} - \cdots + z'_{110} - z'_{111} .$$

Equivalently, c' is a cube if and only if it is obtained by acting on c pointwise by an element of $\mathrm{HK}^3(\mathbb{R}/\mathbb{Z})$, where \mathbb{R}/\mathbb{Z} is given the degree 2 filtration $\mathbb{R}/\mathbb{Z} = \mathbb{R}/\mathbb{Z} = \mathbb{R}/\mathbb{Z} \supseteq \{0\}$.

In summary, for the Heisenberg nilmanifold we have maps

$$\mathcal{H}/\Gamma \xrightarrow{\pi} \mathbb{R}^2/\mathbb{Z}^2 \rightarrow \{*\}$$

where for each map the fibers have a simply transitive action by a compact abelian group (\mathbb{R}/\mathbb{Z} and $\mathbb{R}^2/\mathbb{Z}^2$ respectively); and there is an induced map

$$\mathrm{HK}^k(\mathcal{H}_\bullet)/\Gamma \xrightarrow{\pi} \mathrm{HK}^k(\mathbb{R}^2/\mathbb{Z}^2) \rightarrow \{*\}$$

where the fibers are given by HK^k of the corresponding compact abelian groups, equipped with the “trivial” filtrations of degree 2 and 1 respectively.

It is not hard to argue that a similar property holds for any nilmanifold of any degree, not just the Heisenberg nilmanifold. For a filtered group of degree s , one gets a sequence of maps

$$G \rightarrow G/G_s \rightarrow G/G_{s-1} \rightarrow \cdots \rightarrow G/G_2 \rightarrow G/G_1$$

where typically $G = G_1$ and so the last step is the trivial group. Again, these induce maps of nilmanifolds

$$G/\Gamma \rightarrow (G/G_s)/(\Gamma/(\Gamma \cap G_s)) \rightarrow \cdots$$

where the fibers are all compact abelian groups. A similar compatibility property of the Host–Kra cubes also holds.

In all cases so far, we have made heavy use of the group structure of G and the precise description of Host–Kra cubes. It is a slightly surprising but very important fact that the same conclusions can be made directly in the completely abstract setting of general nilspaces. In other words, we do not need to know about the global group operation on G to recover this tower of compact abelian groups. Indeed, this conclusion holds even for nilspaces that are not of the form G/Γ for any group G .

6.3. The canonical factors. Our first task is to recover a quotient map related to $G \rightarrow G/G_{s+1}$ in the abstract setting. Note that for any filtered group G , the quotient G/G_{s+1} is a filtered group of degree at most s , and therefore has the unique completion property discussed above. It is this that we shall attempt to recover.

Recall Definitions 5.2 and 5.3 (quotient cubespace and the canonical equivalence relation). We first verify the properties of \sim_s that we claimed in Section 5.1. See also [ACS12, Section 2.4] and [HK08, Section 3.3].

Proposition 6.3. *Let $(X, C^k(X))$ be a compact cubespace with k -completion for all k , and let $s \geq 0$ be an integer. Then the canonical equivalence relation \sim_s is indeed an equivalence relation and satisfies the following universal replacement property. If $k \leq s + 1$ and $c \in C^k(X)$, $c' : \{0, 1\}^k \rightarrow X$, $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{1}$ and $c(\vec{1}) \sim_s c'(\vec{1})$, then $c' \in C^k(X)$.*

In particular, X / \sim_s has $(s + 1)$ -uniqueness. Moreover, \sim_s is the smallest equivalence relation with this property. Finally, X / \sim_s still has k -completion for all k , and hence is a nilspace of degree s .

Remark 6.4. Note we are using a convention of stating results using $\vec{1}$ as the “special” vertex of the discrete cube $\{0, 1\}^k$, whenever one is required. However, using the various morphisms of discrete cubes from the definitions of a cubespace, which act transitively on $\{0, 1\}^k$, these properties hold just as well for any fixed vertex in $\{0, 1\}^k$.

First let us see why the definition of \sim_s and the universal replacement property imply $(s + 1)$ -uniqueness of the quotient. Suppose \tilde{c} and \tilde{c}' are two cubes in $C^{s+1}(X / \sim_s)$ such that $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{1}$. By definition of the quotient cubespace, there are cubes $c, c' \in C^{s+1}(X)$ such that $\pi_s(c) = \tilde{c}$, $\pi_s(c') = \tilde{c}'$. But by repeated application of the universal replacement property, *any* configuration $c : \{0, 1\}^{s+1} \rightarrow X$ such that $\pi_s(c) = \tilde{c}$ is a cube, and similarly for c' . Hence we are free to choose c, c' such that $c(\omega) = c'(\omega)$ for each $\omega \neq \vec{1}$.

By definition of \sim_s , we now see that $c(\vec{1}) \sim_s c'(\vec{1})$, and so $\tilde{c}(\vec{1}) = \tilde{c}'(\vec{1})$ as required.

Remark 6.5. Consider briefly the case $s = 0$. Then $x \sim_0 y$ if and only if $[x, y]$ is a 1-cube. It is immediate from the glueing property that this is an equivalence relation. If X is ergodic, then $X / \sim_0 = \{*\}$ is a one-point space. More generally, this identifies the “ergodic components” of X .

It is clear from the hypotheses about the spaces $C^k(X)$ being closed and X being compact that \sim_s is a closed relation.

The following more explicit characterization of \sim_s will be very helpful. See also [ACS12, Lemma 2.3] and [HK08, Proposition 3].


Lemma 6.6. *We have $x \sim_s y$ if and only if the configuration $c : \{0, 1\}^{s+1} \rightarrow X$ given by*

$$c(\omega) = \begin{cases} y & : \omega = \vec{1} \\ x & : \omega \neq \vec{1} \end{cases}$$

is an $(s + 1)$ -cube.

Proof sketch. The “if” direction is straightforward, since the constant configuration $c'(\omega) = x$ for all ω is in $C^{s+1}(X)$ by the axioms (see Remark 3.3). We illustrate the “only if” direction with a picture when $s = 1$; the general case is similar but notationally awkward.

We know there exist c, c' with $c(\omega) = c'(\omega)$ for $\omega \neq \vec{1}$, $c(\vec{1}) = x$ and $c'(\vec{1}) = y$. Consider the picture




It is clear each of the small cubes is a cube, as they are just rotated and reflected copies of c or c' . By two applications of the glueing property, the outer square is also a cube, as required. \square


We can now check:

Proposition 6.7. *The relation \sim_s is an equivalence relation.*

Proof sketch. Symmetry and reflexivity are immediate from the original definition. For transitivity, we apply the lemma; again we will give a pictorial sketch when $s = 1$. Suppose $x \sim_s y$, $y \sim_s z$ and consider




where the left and right squares are cubes by the lemma and a reflection. Then the outer rectangle is a square by the glueing property. But so is (say)



by a duplication operation (i.e. using the cubespace axioms), and so $x \sim_s z$ from the definition, as required. \square

Sketch proof of Proposition 6.3. We prove the universal replacement property. First suppose $k = s + 1$. Take c, c' as in the statement, and write $x = c(\vec{1})$, $y = c'(\vec{1})$. Again we give a picture:



The bottom left square is just c , the top right is a cube by the lemma, and the other two small squares are cubes again by a duplication. Hence the outer square is a cube by glueing, as required.

Now suppose $k < s + 1$. Given a cube $c \in C^k(X)$, using appropriate morphisms of discrete cubes and the cubespace axioms, we may

- duplicate c up to an $(s + 1)$ cube $\tilde{c} \in C^{s+1}(X)$;
- change entries of \tilde{c} repeatedly using the above; and
- restrict to some appropriate face of \tilde{c} to obtain a cube of dimensions k with the desired properties.

The statement about the completion property for X/\sim_s is deferred to Section 7. \square

Note it is clear that X/\sim_s is also compact, and ergodic if X is.

Remark 6.8. We observe that this quotient π_s as constructed is completely canonical: firstly in that it depends only on the cubespace $(X, C^k(X))$, and secondly that it is “minimal” for the properties required of it in some sense.

One way of stating the latter is that $\pi_s(X)$ obeys a certain universal property: if $(Y, C^k(Y))$ is any other nilspace of degree s then any cubespace morphism $X \rightarrow Y$ factors through $\pi(X)$.

In particular, \sim_s is trivial if and only if X has $(s + 1)$ -uniqueness.

The full strength of these canonical factors comes when they are chained together. That is, as before we can construct a tower of maps

$$X \xrightarrow{\pi_s} \pi_s(X) \xrightarrow{\pi_{s-1}} \pi_{s-1}(X) \rightarrow \dots \xrightarrow{\pi_0} \pi_0(X)$$

where $\pi_0(X) = \{*\}$, the one-point space, provided X is ergodic. To check this makes sense, we need to verify the following trivial observation.

Proposition 6.9. *The equivalence relations \sim_s are nested, i.e. if $t \geq s$ and $x \sim_t y$ then $x \sim_s y$. Moreover, $\pi_s(\pi_t(X)) = \pi_s(X)$, i.e. the definition of \sim_s is not affected by first quotienting by \sim_t .*

Proof. If $c, c' \in C^{t+1}(X)$ are any cubes verifying $x \sim_t y$ then any sub-cube of dimension $(s+1)$ containing $\vec{1}$ will verify $x \sim_s y$. \square

6.4. Structure groups and the weak structure theorem. Having established the existence of the tower

$$\dots \xrightarrow{\pi_s} \pi_s(X) \xrightarrow{\pi_{s-1}} \pi_{s-1}(X) \rightarrow \dots \xrightarrow{\pi_0} \pi_0(X)$$

in the abstract, to further the analogy with the Heisenberg case (or more generally, the case of any nilmanifold) we need to describe the *fibers* of each map π_t . Recall these are expected to have the structure of a compact abelian group; or more precisely, to have a free and transitive action by a compact abelian group.

As in Section 5.1 we let $\mathcal{D}_s(A)$ denote the Host–Kra cubespace on an Abelian group equipped with the filtration $A = A_0 = A_1 = \dots = A_s \supseteq \{0\}$, which is a nilspace of degree s . Recall also the statement of the weak structure theorem (Theorem 5.4).

Remark 6.10. Suppose X is an ergodic nilspace of degree 1, and consider the weak structure theorem in the case $s = 1$. Since $\pi_0(X) = \{*\}$ (see Remark 6.5), the theorem asserts precisely that X is isomorphic to $\mathcal{D}_1(A)$ for some compact abelian group A , i.e. X bijects with A and the cubes of X are identified with the Host–Kra cubes of A with the degree 1 filtration.

(This isomorphism is found by fixing $x_0 \in X$ and identifying $A \leftrightarrow X$ by $a \leftrightarrow a(x_0)$. In particular this is canonical only up to the choice of x_0 .)

A proof of this case is sketched above in Section 4.2. However, that proof does not generalize entirely cleanly to the case $s > 1$. We will now sketch a slightly different argument that does generalize, but working again in the case $s = 1$ for ease of notation.


First we will sketch how the proof would look if we already knew that $X = \mathcal{D}_1(A)$ for some A . We consider all the edges ⁶ $[a, b]$ of X , and associate to each one the group element $(b - a)$. We define an equivalence relation on edges by $[a, b] \sim [a', b']$ if $b - a = b' - a'$; so A is precisely the set of equivalence classes of edges under this relation.

Given edges $[a, b]$ and $[b, c]$ we can concatenate them to get $[a, c]$. If the associated elements of A are $r = b - a$ and $s = c - b$ then $[a, c]$ is associated to $c - a = r + s$. Thus we can recover the group addition operation on A in a combinatorial fashion by concatenation. The other operations are similarly easy to define.

The task is now to show that these same constructions make sense without any *a priori* assumptions on X other than that it is a nilspace of degree 1.

Sketch proof of Theorem 5.4, $s = 1$. Since X is ergodic, the edges $C^1(X)$ are in bijection with $X \times X$. For two edges $[x, y]$ and $[x', y']$ we write $[x, y] \sim [x', y']$ whenever


⁶An edge is just another name for a 1-cube.



is a 2-cube. This is an equivalence relation by the cubespace axioms and glueing. As a set, we define $A := (X \times X) / \sim$.

Claim 6.11. *Let $[x, y]$ be an edge. For any fixed $e \in X$ there is an unique $a \in X$ such that $[x, y] \sim [e, a]$. In other words, the class of $[x, y]$ in A has an unique representative of the form $[e, a]$.*

Proof of claim. This follows trivially from unique completion of the corner




giving a in the top right. □

We now define addition on A by concatenation as suggested above. More precisely, given $g, h \in A$, we fix an $e \in X$ and choose the unique representatives $g \sim [x, e]$ and $h \sim [e, y]$; then $g + h$ is defined to be the class of the concatenation $[x, y]$.

Claim 6.12. *This definition of $g + h$ is well-defined; i.e. the outcome does not depend on the choice of e .*

Proof of claim. Considering the diagram




where $[x, e] \sim [x', e'] \sim g$ and $[e, y] \sim [e', y'] \sim h$, since the outer square is a 2-cube by glueing we have $[x, y] \sim [x', y']$ as required. □

The identity $0 \in A$ is given by the class of constant edges $[x, x]$ (by 2-uniqueness this is indeed a class of \sim), and inversion by the operation $[x, y] \mapsto [y, x]$.

Claim 6.13. *The set A with these operations forms a (topological) abelian group.*


Proof of claim. That the proposed identity and inverse operation do what they claim is a trivial check. Associativity follows from associativity of concatenation, together with Claim 6.12. For commutativity, suppose $[x, e] \sim g$, $[e, y] \sim h$; by completing the corner to get e' in



we have $[e', y] \sim g$, $[x, e'] \sim h$ and hence $h + g \sim [x, y] \sim g + h$ as required.


Issues of continuity are not difficult to justify, and we will not do so here. \square

We now define the group action of A on X . Given $g \in A$ and $x \in X$ we take $g(x)$ to be the unique element of X such that $g \sim [x, g(x)]$. It is clear $0 \in A$ acts trivially; looking at the diagram




where $g \sim [c, e]$ and $h \sim [e, d]$, we conclude this is indeed a group action by considering the outer square. Moreover, it is trivially simply transitive.

Finally we must investigate the cubes of X in terms of A . For 0 and 1-cubes there is nothing to say. Suppose c is the configuration



in $X^{\{0,1\}^2}$. By definition of \sim , this is a cube if and only if $[x, y] \sim [z, w]$. Furthermore it is obtained by acting on the constant cube



by the elements of A represented by $r := [e, x]$, $s := [e, y]$, $t := [e, z]$, $u := [e, w]$ respectively.

By considering the concatenation of $[x, e]$, $[e, y]$ and using the definitions of the group operation on A , we find that $[x, y] \sim s - r$, and similarly $[z, w] \sim u - t$. Hence, c is a cube if and only if $[x, y] \sim [z, w]$, if and only if $s - r = u - t$; i.e. if and only if c is obtained by acting on a constant cube by an element of $C^2(\mathcal{D}_1(A))$, as required.

It is straightforward to deduce the cases $\ell \geq 3$ from this, but we defer this argument to the general case. \square

Several adjustments are required to make this argument work in the full case of Theorem 5.4. However, we will address these at the same time as stating and proving a version of Theorem 5.4 in greater generality.

Although the statement of Theorem 5.4 refers to only map $\pi_{s-1}: X \rightarrow \pi_{s-1}(X)$, we can exploit the fact that $\pi_{s-1}(X)$ is itself a nilspace of degree $(s-1)$ to apply the theorem repeatedly, and obtain a tower

$$X \xrightarrow{\pi_{s-1}} \pi_{s-1}(X) \rightarrow \dots \xrightarrow{\pi_0} \pi_0(X) = \{*\}$$

as promised. We write A_t (for $1 \leq t \leq s$, if X is a nilspace of degree s) for the compact abelian groups that arise from the application of Theorem 5.4 to $\pi_t(X)$ at each stage; collectively, these are referred to as the *structure groups* of X .

7. THE RELATIVE WEAK STRUCTURE THEORY

With these overviews concluded, we will now return to a formal account of the weak structure theory. In this second pass, we will both give complete proofs, and also introduce a slightly greater degree of generality which will be useful to us in future.

7.1. Fibrations. One key definition missing from our initial treatment of cubespaces and nilspaces is the following notion of a *fibration*⁷, which is a particular kind of cubespace morphism. This can be thought of as a relative version of the notion of the corner completion property. An alternative heuristic is that fibrations are morphisms which are “properly” surjective in the category of cubespaces.

Definition 7.1. Let $f: X \rightarrow Y$ be a morphism between cubespaces $(X, C^k(X))$ and $(Y, C^k(Y))$. We say f is a *fibration* if the following holds for any $k \geq 0$.

Suppose $c \in C^k(Y)$ is a k -cube, and $\lambda: \{0, 1\}^k \setminus \{\vec{1}\} \rightarrow X$ is a k -corner (see Definition 3.5) such that $f(\lambda(\omega)) = c(\omega)$ for all $\omega \in \{0, 1\}^k \setminus \{\vec{1}\}$. Then there exists $x \in X$ such that

- the configuration

$$\begin{aligned} \tilde{c}: \{0, 1\}^k &\rightarrow X \\ \omega &\mapsto \begin{cases} \lambda(\omega) & : \omega \neq \vec{1} \\ x & : \omega = \vec{1} \end{cases} \end{aligned}$$

⁷This term was chosen by analogy with the notion of a Kan fibration in simplicial homotopy theory. We believe this and other resemblances between these two theories are not entirely superficial, but we will not pursue this here.

- is in $C^k(X)$; and
- $f(x) = c(\tilde{1})$ (and hence $f(\tilde{c}) = c$).

Informally, this somewhat hard-to-parse definition says: given any k -corner λ of X and any completion of $f(\lambda)$ to a cube of Y , we can complete λ to a cube in a *compatible* fashion.

Remark 7.2. By taking $k = 0$, we find that any fibration is surjective as a map $X \rightarrow Y$. However, this is strictly weaker, as can be seen by taking $X = \mathcal{D}_1(\mathbb{R}/\mathbb{Z})$, $Y = \mathcal{D}_2(\mathbb{R}/\mathbb{Z})$ and $f: \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ the identity map. This is a surjective cubespace morphism, but taking $k = 2$ and choosing c to be any configuration in $(\mathbb{R}/\mathbb{Z})^4$ that is not a cube of $C^2(X)$ (but is in $C^2(Y) = (\mathbb{R}/\mathbb{Z})^4$) it is seen not to be a fibration.

Remark 7.3. Let X be a cubespace and consider the unique morphism $f: X \rightarrow \{*\}$, the one-point cubespace. Then f is a fibration if and only if X has k -completion for all k .

Hence, we refer to spaces that have k -completion for all k as *fibrant*.

By the previous remark, whenever we prove a statement about fibrant cubespaces, it is reasonable – and sometimes useful – to ask for a relative version that holds for fibrations.

Remark 7.4. The notion of a fibration is intimately related to that of a *fiber-surjective morphism* which appears in [ACS12, Section 2.8]. There, a cubespace morphism $f: X \rightarrow Y$ is called fiber-surjective if for each $k \geq 0$, the image of any \sim_k class in X is a \sim_k class of Y . It is assumed in this context that both X and Y are nilspaces.

It is not hard to check that if X and Y are nilspaces and $f: X \rightarrow Y$ is a cubespace morphism, then f is a fibration if and only if it is fiber-surjective. If X and Y are general cubespaces, it is not clear how to extend the definition of fiber-surjectivity in general, so no comparison is possible. If one did extend it to a wider class of cubespaces, we note that the map $X \rightarrow \{*\}$ will surely be vacuously fiber-surjective for any X , whereas it is a fibration if and only if X is fibrant, and hence the notions would be inequivalent whenever X is not assumed to be fibrant.

Our reasons for working with the definition of a fibration given here rather than with fiber-surjective morphisms are twofold.

- We will have cause later in the project to work with fibrations $f: X \rightarrow Y$ where X and Y are not nilspaces. Here the distinction matters, and the notion of a fibration is the correct one for what we need.
- The fact that the definition of a fibration cleanly extends the completion axiom means that many proofs that work for nilspaces tend to transfer without significant modification to relative versions for fibrations. In the authors' view, this makes the theory cleaner in places, and suggests the fibration definition is perhaps the more natural one.

However, we stress that the difference is a technical one, and logically unimportant in almost all cases.

We will now show a couple of lemmas, together with their proofs, to illustrate this concept in action.

Lemma 7.5 (Lifting of “partial cubes”). *Let $f: X \rightarrow Y$ be a fibration between cubespaces X and Y .*

Let $S \subseteq \{0, 1\}^k$ be a downwards-closed subset, i.e. if $\omega \in S$ and $\omega' \subseteq \omega$ then $\omega' \in S$. Let T be another such subset with $S \subseteq T$.

Suppose we have configurations $A: S \rightarrow X$, $B: T \rightarrow Y$ such that

- (i) for each $\omega \in S$, we have $f(A(\omega)) = B(\omega)$;*
- (ii) for each $\omega \in S$, the configuration*

$$\begin{aligned} \{\omega': \omega' \subseteq \omega\} &\rightarrow X \\ \omega' &\mapsto S(\omega') \end{aligned}$$

is a cube of X ; and

- (iii) similarly for Y and T .*

Then we can extend A to a function $\tilde{A}: T \rightarrow X$ such that $f(\tilde{A}) = B$ and \tilde{A} now has property (ii) with respect to T .

Here we again used the identification of $\{0, 1\}^k$ with the set of subsets of $[k]$ as in Section 2. This rather general statement has the following more natural corollaries.

Corollary 7.6. *If $f: X \rightarrow Y$ is a fibration then the induced map $f: C^k(X) \rightarrow C^k(Y)$ is surjective for all $k \geq 0$.*


Proof. Set $S = \emptyset$, $T = \{0, 1\}^k$ in the lemma. □

Corollary 7.7. *If X is fibrant, $S \subseteq \{0, 1\}^k$ is downward closed and $A: S \rightarrow X$ is a configuration satisfying property (ii) from the lemma, then A extends to a cube, i.e. there is some cube $c \in C^k(X)$ with $c|_S = A$.*

Proof. Apply the lemma with $Y = \{*\}$, $T = \{0, 1\}^k$. □

Proof of Lemma 7.5. If $T = S$ there is nothing to do. Suppose ω_0 is a minimal element of $T \setminus S$. If we can extend A to a configuration on $S \cup \{\omega_0\}$ that still satisfies (i) and (ii), we will be done by iterating this process.

An example of this set-up is shown in the following diagram, where $T = \{0, 1\}^3 \setminus \{\vec{1}\}$, $S = \{000, 001, 010, 011, 100\}$ and $\omega_0 = 101$:




Let $V = \{\omega \in \{0, 1\}^k : \omega \subseteq \omega_0\}$; clearly this is a discrete sub-cube of dimension $k' = |\omega_0|$ (shown as the dotted square in the diagram, where $k' = 2$). Then $A|_{V \setminus \{\omega_0\}}$ is a k' -corner, and furthermore a partial lift of the corresponding cube $B|_V$ of Y .

By the fibration definition we can choose $x \in X$ such that $f(x) = B(\omega_0)$, and such that setting $A(\omega_0) = x$ we have $A|_V \in C^{k'}(X)$. Hence this extended A satisfies (i) and (ii) as required. \square

We can also use this to give a full proof of Proposition 6.2; recall this stated that a fibrant space satisfies the glueing axiom.

Proof of Proposition 6.2. Essentially we formalize the diagram given above when $k = 2$; for convenience we reproduce it now.



Formally: given $c, c' \in C^k(X)$ with $c(\omega 1) = c'(\omega 0)$ for all $\omega \in \{0, 1\}^{k-1}$, let $S = \{0, 1\}^{k+1} \setminus (\{0, 1\}^{k-1} \times \{(1, 1)\})$ and consider the configuration

$$\begin{aligned} d: S &\rightarrow X \\ \omega 00 &\mapsto c(\omega 1) = c'(\omega 0) \\ \omega 01 &\mapsto c'(\omega 1) \\ \omega 10 &\mapsto c(\omega 0) . \end{aligned}$$

One can verify this satisfies the hypotheses of Corollary 7.7, and so there is a cube $\tilde{d} \in C^{k+1}(X)$ whose restriction to S is d . But then the “glued” configuration c'' as in Definition 6.1 is a sub-cube of \tilde{d} as shown, and so is in $C^k(X)$ by the cubespace axioms. \square

It is clear from the definition that composition of fibrations is a fibration. Another useful fact is the following universal property.

Lemma 7.8 (“Universal property”). *Let $f_{YX} : X \rightarrow Y$ and $f_{ZX} : X \rightarrow Z$ be fibrations between compact cubespaces. Suppose that for every $y \in Y$ there is $z \in Z$ such that $f_{YX}^{-1}(y) \subseteq f_{ZX}^{-1}(z)$. Then there is a unique fibration $f_{ZY} : Y \rightarrow Z$ such that $f_{ZX} = f_{ZY} \circ f_{YX}$.*

Equivalently, the following holds. Let $f : X \rightarrow Y$ be a fibration and $g : Y \rightarrow Z$ be a map between two compact cubespaces. If $g \circ f$ is a fibration then so is g .

Proof. Since X , Y and Z are compact metric spaces, f_{YX} and f_{ZX} are quotient maps. Hence the map f_{ZY} (which is uniquely defined thanks to the condition imposed on the fibres of f_{YX} and f_{ZX}) is continuous. It remains to show that it is also a fibration, which is precisely the second statement.

For the second statement, we first show that g is a cubespace morphism. To this end, fix a cube $c \in C^k(Y)$ and let $\tilde{c} \in C^k(X)$ be a cube such that $f(\tilde{c}) = c$. This exists since f is a fibration. This implies that $g(c) = g \circ f(\tilde{c}) \in C^k(Z)$, hence g is indeed a cubespace morphism.

Suppose λ is a k -corner in Y and $c \in C^k(Z)$ a compatible k -cube. By Lemma 7.5 on f with $S = \emptyset$, $T = \{0, 1\}^k \setminus \{\vec{1}\}$, we can choose $\tilde{\lambda}$ a k -corner of X such that $f(\tilde{\lambda}) = \lambda$ and so in particular $g \circ f(\tilde{\lambda})$ is compatible with c .

Since $g \circ f$ is a fibration, we may extend $\tilde{\lambda}$ to a cube $c' \in C^k(X)$ such that $g \circ f(c') = c$. So, $f(c') \in C^k(Y)$ has the required property. \square

Remark 7.9. Taking $Z = \{*\}$ in the lemma, we see that if $f : X \rightarrow Y$ is a fibration and X is fibrant then so is Y .

It also holds that the property of s -uniqueness is inherited by the image of a cubespace under a fibration, hence the image of a nilspace under a fibration is a nilspace.

Indeed, let $f : X \rightarrow Y$ be a fibration and suppose that X has s -uniqueness. Suppose that $x \sim_s y$ for some $x, y \in Y$. We aim to show that $x = y$. We first note that the configuration $c : \{0, 1\}^s \rightarrow Y$

defined by $\omega \mapsto x$ for $\omega \neq \vec{1}$ and $\vec{1} \mapsto y$ is a cube in $C^s(Y)$. Let $\tilde{x} \in f^{-1}(x)$ be arbitrary and let $\lambda : \{0, 1\}^s \setminus \{\vec{1}\} \rightarrow X$ be the constant corner configuration $\omega \mapsto \tilde{x}$.

Since f is a fibration, λ can be completed to a cube \tilde{c} such that $f(\tilde{c}(\vec{1})) = y$. Since X has s -uniqueness, we have $\tilde{x} = \tilde{c}(\vec{1})$, which implies $x = f(\tilde{x}) = f(\tilde{c}(\vec{1})) = y$, as required.

We also need to record the correct “relative” version of uniqueness.

Definition 7.10. We say a cubespace morphism $f : X \rightarrow Y$ has k -uniqueness if the following holds: if $c, c' \in C^k(X)$ are two cubes such that $f(c) = f(c')$ and $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{1}$, then in fact $c = c'$.

If f is a fibration and k is the smallest number such that f has k -uniqueness, we say that f has *degree* $(k - 1)$.

Again it is clear that a space X has k -uniqueness if and only if the map $X \rightarrow \{*\}$ does.

7.2. The structure theory in terms of fibrations. We will now modify – and prove – statements of the weak structure theory discussed above, in “relative form”, i.e. in terms of general fibrations.

First we consider the *canonical factors* π_s . Everything that happens here should be a reasonably straightforward generalization of what happened in Section 6.3.

Definition 7.11. Let $f : X \rightarrow Y$ be a fibration, and let $k \geq 0$. Define an equivalence relation $\sim_{f,s}$ on X as follows: $x \sim_{f,s} x'$ if there exist two $(s + 1)$ -cubes c, c' in X such that $f(c) = f(c')$, $c(\vec{1}) = x$, $c'(\vec{1}) = x'$ and $c(\omega) = c'(\omega)$ at all other vertices $\omega \in \{0, 1\}^{s+1} \setminus \{\vec{1}\}$.

Proposition 7.12. Suppose $f : X \rightarrow Y$ is a fibration between compact cubespaces X, Y that have the glueing property.

Then the relation $\sim_{f,s}$ is a closed equivalence relation. Moreover a “universal replacement” property holds: if $x \sim_{f,s} x'$, $k \leq s + 1$ and $c \in C^k(X)$ is a cube with $c(\vec{1}) = x$, then the configuration c' given by

$$c'(\omega) = \begin{cases} x' & : \omega = \vec{1} \\ c(\omega) & : \omega \neq \vec{1} \end{cases}$$

is a cube.

Writing $\pi = \pi_{f,s} : X \rightarrow X / \sim_{f,s}$ for the projection map, we have that π is a fibration and f factors as $f : X \xrightarrow{\pi} X / \sim_{f,s} \xrightarrow{g} Y$ where g is a fibration of degree at most s .

The reader may have observed that $x \sim_{f,s} y$ precisely if $x \sim_s y$ and $f(x) = f(y)$, and therefore many (but not all) of the above properties following immediately from the corresponding properties of \sim_s . However, since we still owe the reader a complete proof even in that non-relative case, we will lose nothing by starting from scratch in full generality. However, the reader may wish to refer to Section 6.3 for a conceptually similar but notationally less painful discussion.

We will first introduce some symbology for manipulating high-dimensional cubes.

Definition 7.13. If c and c' are k -cubes, we will denote by $[c, c']$ the $(k+1)$ -configuration


$$\omega \mapsto \begin{cases} c(\omega_1 \dots \omega_k) & : \omega_{k+1} = 0 \\ c'(\omega_1 \dots \omega_k) & : \omega_{k+1} = 1 . \end{cases}$$

Given an element $x \in X$, the notation $\square^k(x)$ denotes the constant k -cube ($\omega \mapsto x$). Given $x, y \in X$, we denote by $\sqcup^k(x; y)$ the configuration


$$\{0, 1\}^k \rightarrow X$$

$$\omega \mapsto \begin{cases} x & : \omega \neq \vec{1} \\ y & : \omega = \vec{1} . \end{cases}$$

We reserve the right to combine these pieces of notation freely with each other and also with pictorial representations of cubes. For instance, the notation $[\sqcup^2(x; y), [\square^1(z), \square^1(w)]]$ is a synonym for



and



although we will rarely need to denote anything quite so unpleasant.


We will be generally cavalier about the ordering of the indices $\{1, \dots, k\}$ implied by successive applications of this notation, since it should always be clear from context what is meant, and because (thanks to the axioms) it makes little or no difference most of the time.

Finally, we will need to introduce some notation about *tri-cubes* of the form appearing in the proof of Proposition 6.3, i.e. a collection of 2^k k -cubes that all glue together at the middle. These are very useful gadgets introduced by Antolín Camarena and Szegedy in [ACS12, Section 2.3].

Definition 7.14. We say a collection $t = (t_\nu)_{\nu \in \{0,1\}^k}$ constitutes a *tri-cube* if $t_\nu \in C^k(X)$ for each ν , and if for each $\nu, \nu', \omega \in \{0,1\}^k$ such that for each $i \in \{1, \dots, k\}$ either $\nu_i = \nu'_i$ or $\omega_i = 1$, we have $t_\nu(\omega) = t_{\nu'}(\omega)$.

We say the *outer cube* of t is the configuration $\omega \mapsto (t_\omega(\vec{0}))$.

Note this corresponds to a picture (when $k = 2$):



and hence some of the t_ν are reflected relative to the standard orientation.

Proposition 7.15. *If X is a cubespace with the glueing property and (t_ν) is a tri-cube in X , then the outer cube is in $C^k(X)$.*

Proof. This follows from repeated application of the glueing property. By assumption we can glue together each pair $t_{\eta 0}$ and $t_{\eta 1}$ for $\eta \in \{0,1\}^{k-1}$, and then repeat for each of the other $(k-1)$ coordinates in turn. \square

Proof of Proposition 7.12. The proof is strongly analogous to that from Section 6.3. We first prove an analogue of Lemma 6.6.

Lemma 7.16. *We have that $x \sim_{f,s} y$ if and only if $f(x) = f(y)$ and also $\perp^{s+1}(x; y) \in C^{s+1}(X)$.*

Proof. The strategy is to construct a tri-cube whose outer cube is $\perp^{s+1}(x; y)$.

By assumption we are given cubes $c, c' \in C^{s+1}(X)$ such that $f(c) = f(c') = d$, $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{0}$ and $c(\vec{0}) = x$, $c'(\vec{0}) = y$. (Note we have switched $\vec{1}$ for $\vec{0}$ for convenience.) From this it is immediate that $f(x) = f(y)$.

Now define the family $(t_\nu)_{\nu \in \{0,1\}^{s+1}}$ by

$$t_\nu = \begin{cases} c & : \nu \neq \vec{1} \\ c' & : \nu = \vec{1} . \end{cases}$$

It is clear from the hypotheses that this is a tri-cube with outer cube $\perp^{s+1}(x; y)$ as required. \square

Claim. *The relation $\sim_{f,s}$ is an equivalence relation.*

Proof of claim. Again, symmetry and reflexivity are immediate. For transitivity, assume $x \sim_{f,s} y \sim_{f,s} z$. Then by the lemma, $f(x) = f(y) = f(z)$ and $\sqcup^{s+1}(y; x), \sqcup^{s+1}(y; z)$ are cubes. By glueing along the common face $\square^s(y)$, we find that $\sqcup^s(\square^1(y); [x, z])$ is an $(s+1)$ -cube. But so is $\sqcup^s(\square^1(y); [x, x])$, by restricting to a face of $\sqcup^{s+1}(y; x)$ and then duplicating. Since these latter two cubes have the same image under f , we have $x \sim_{f,s} z$ by the definition. \square

Again, the fact that this relation is closed is straightforward from the definition, the hypothesis that $C^{s+1}(X)$ is closed, continuity of f and compactness.

We turn to the “universal replacement” statement.

Claim. *If $x \sim_{f,s} x'$, and $c \in C^{s+1}(X)$ is a cube with $c(\vec{1}) = x$, then the configuration c' given by*

$$c'(\omega) = \begin{cases} x' & : \omega = \vec{1} \\ c(\omega) & : \omega \neq \vec{1} \end{cases}$$

is a cube.

As in the proof of Proposition 6.3, this easily implies the same statement for all $k < s+1$.

Proof of claim. We construct a tri-cube whose outer cube is c' . We set

$$t_\nu(\omega) = \begin{cases} c((\min(\nu_i + \omega_i, 1))_{i=1}^{s+1}) & : \nu \neq \vec{1} \\ x & : \nu = \vec{1}, \omega \neq \vec{0} \\ x' & : \nu = \vec{1}, \omega = \vec{0} \end{cases}$$

which is a very elaborate way of stating the generalization of the diagram from the proof of Proposition 6.3. Now, $t_{\vec{0}}$ is just c , the other t_ν except $t_{\vec{1}}$ are obtained from c by duplication of an upper face, and $t_{\vec{1}}$ is a reflected copy of $\sqcup^{s+1}(x; x')$. So, all the t_ν are cubes, and a tedious check verifies that (t_ν) is a tri-cube with outer cube c' as required. \square

So, we have a projection $\pi_{f,s}: X \rightarrow X / \sim_{f,s}$, and by construction f factors through $\pi_{f,s}$, i.e. $f = g \circ \pi_{f,s}$ for some well-defined function $g: X / \sim_{f,s} \rightarrow Y$.

Claim. *The map $\pi_{f,s}$ is a fibration.*

Proof of claim. Suppose we are given a k -corner λ in X and a compatible $c \in C^k(X) / \sim_{f,s}$.

First consider the case $k \leq s+1$. By definition there is a cube \tilde{c} of X with $\pi_{f,s}(\tilde{c}) = c$, and furthermore by repeated application the second claim, *any* configuration $\tilde{c}: \{0, 1\}^k \rightarrow X$ with $\pi_{f,s} \circ \tilde{c} = c$ is a cube. Hence taking $\lambda(\vec{1})$ to be an arbitrary point of $\pi_{f,s}^{-1}(c(\vec{1}))$ must work.

If $k \geq s+1$, we may use the fact that f is a fibration to complete λ to a cube \tilde{c} such that $f(\tilde{c}) = g(c)$. But now $\pi_{f,s}(\tilde{c})$ and c are two cubes of $X / \sim_{f,s}$ whose image under g is the same, and which differ

only at one vertex. Passing to an upper face of dimension $(s + 1)$ if necessary and applying the third claim, we find $\pi_{f,s}(\tilde{c}) = c$ as required. \square

We now note that g is therefore a fibration by the universal property (Lemma 7.8).

Claim. *The map g has $(s + 1)$ -uniqueness.*

Proof of claim. Indeed, given two $(s + 1)$ -cubes c, c' of $X / \sim_{f,s}$ with $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{1}$ and $g(c) = g(c')$, we know there are cubes \tilde{c}, \tilde{c}' of X such that $\pi_{f,s}(\tilde{c}) = c$, $\pi_{f,s}(\tilde{c}') = c'$. By repeated application of the second claim we may assume $\tilde{c}(\omega) = \tilde{c}'(\omega)$ for each $\omega \neq \vec{1}$. But then $\tilde{c}(\vec{1}) \sim_{f,s} \tilde{c}'(\vec{1})$ by the definition of $\sim_{f,s}$, and so $c = c'$. \square

This completes the proof of Proposition 7.12. \square

Remark 7.17. Again we can consider the case that $Y = \{*\}$ and X is a fibrant cubespace. In this case, Proposition 7.12 essentially restates Proposition 6.3. Note in particular that we have now proven that X / \sim_s is fibrant (since $g: X / \sim_s \rightarrow \{*\}$ is a fibration).

7.3. The (relative) structure groups. With this in place, we are in a position to attack the relative analogue of Theorem 5.4. First, we need one further “relative” definition.

Definition 7.18. A morphism $f: X \rightarrow Y$ is called *(relatively) k -ergodic* if for any $1 \leq \ell \leq k$ and $c \in C^\ell(Y)$, any configuration in $f^{-1}(c)$ is a cube of X .

We say it is *relatively ergodic* if it is relatively 1-ergodic.

The direct analogue of Theorem 5.4 is as follows.

Theorem 7.19. *Suppose X and Y are compact ergodic cubespaces that obey the glueing condition. Let $f: X \rightarrow Y$ be a relatively s -ergodic fibration of degree at most s . Then there exists a compact Abelian group $A = A(f)$ acting continuously on X such that:*

- (i) *the action of A on X is free, and its orbits are precisely the fibers of f ;*
- (ii) *this induces a (free) pointwise action of $C^k(\mathcal{D}_s(A))$ on $C^k(X)$, whose orbits are precisely the fibers of the map $f: C^k(X) \rightarrow C^k(Y)$.*

We will see that the proof is constructive and A is defined canonically in terms of the cube structures and f .

The key point about these hypotheses is that they hold whenever f is the canonical projection map $\pi_{g,s-1}: \pi_{g,s}(X) \rightarrow \pi_{g,s-1}(X)$ defined with respect to some other fibration $g: X \rightarrow Y$. In this case, we write $A = A_s(g)$, the “ s -th structure group of the fibration”. In this way we can iterate this theorem to obtain:

Corollary 7.20. *Let $f: X \rightarrow Y$ be a fibration of degree at most s between compact ergodic cubespaces X, Y that obey the glueing condition. Then we have a tower*

$$X = \pi_{f,s}(X) \xrightarrow{\pi_{f,s-1}} \pi_{f,s-1}(X) \rightarrow \dots \xrightarrow{\pi_{f,0}} \pi_{f,0}(X) = Y$$

where the fibers of $\pi_{f,i-1}$ are identified with the compact abelian group $A_i(f)$ for each i , in the sense of Theorem 7.19.

Proof of Corollary 7.20 from Theorem 7.19. It suffices to know that each map $\pi_{f,t}: \pi_{f,t+1}(X) \rightarrow \pi_{f,t}(X)$ is well-defined, a fibration, has $(t+2)$ -uniqueness and is $(t+1)$ -ergodic.

Well-definedness follows from the nested nature of the relations $\sim_{f,t}$; the proof of this is unchanged from Proposition 6.9.

The fact that it is a fibration follows from the fact that $\pi_{f,t+1}: X \rightarrow \pi_{f,t+1}(X)$ and $\pi_{f,t}: X \rightarrow \pi_{f,t}(X)$ are fibrations, and the universal property (Lemma 7.8).

The $(t+2)$ -uniqueness statement is supplied by Proposition 7.12 applied to $\pi_{f,t+1}(X)$.

Finally, the $(t+1)$ -ergodicity statement is precisely what is encoded by the “universal replacement” property from that same proposition as long as $t \geq 1$. For $t = 0$ it does not follow from Proposition 7.12 (indeed in general it does not hold). Instead for $t = 0$ we use that assumption that X is ergodic. \square

Before embarking on the proof of Theorem 7.19, we note a relative analogue of Proposition 3.11 that will be useful for discussing high-dimensional cubes.

Proposition 7.21. *Let $s \geq 1$ be an integer, $f: X \rightarrow Y$ a fibration of degree at most s and $k \geq s+1$. Further let $c: X \rightarrow \{0,1\}^k$ be given. Then $c \in C^k(X)$ if and only if $f(c)$ is a cube of Y and further every face of c of dimension $(s+1)$ is a cube of X .*

Proof. The “only if” part is clear. For the converse, again we argue by induction on k , the case $k = s+1$ being trivial. Given c , we consider the restriction to $\{0,1\}^k \setminus \{\vec{1}\}$, which by inductive hypothesis is a k -corner. Hence we may complete this corner relative to $f(c)$; i.e. there is a cube c' such that $c'(\omega) = c(\omega)$ for all $\omega \neq \vec{1}$ and $f(c'(\vec{1})) = f(c(\vec{1}))$. Now by restricting to any upper face of c and c' of dimension $(s+1)$ and invoking $(s+1)$ -uniqueness, we find $c(\vec{1}) = c'(\vec{1})$ also and hence $c = c'$ is a cube. \square

The remainder of this section is spent proving Theorem 7.19.

Proof of Theorem 7.19. Again, this account is very similar to the proof of Theorem 5.4, and the reader will find it significantly more comprehensible if they are already familiar with the sketch proof of that theorem given above in Section 6.4.

Before, we defined the structure group in terms of equivalence classes of edges in $C^1(X)$, under the equivalence relation $[x, y] \sim [z, w]$ if and only if $[[x, y], [z, w]] \in C^2(X)$.

One way to rationalize this construction is in terms of the *edge cubespace* $\mathcal{E}(X)$, which we define now.

Definition 7.22. Let X be a cubespace. The *edge cubespace* of X , denoted $\mathcal{E}(X)$, is the cubespace whose base space consists of the edges $C^1(X) \subseteq X \times X$, and whose k -cubes are the configurations $c \in (X \times X)^{\{0,1\}^k} \cong X^{\{0,1\}^{k+1}}$ that correspond to $(k+1)$ -cubes of X .

It is easy to check that this does indeed define a cubespace.

It follows that $[x, y] \sim [z, w]$ as above if and only if $[[x, y], [z, w]]$ is an edge of $\mathcal{E}(X)$, hence if and only if $[x, y] \sim_0 [z, w]$ where \sim_0 denotes the canonical equivalence relation on $\mathcal{E}(X)$.

In order to generalize the proof of Theorem 5.4 to $s > 1$, it is therefore tempting to try replacing \sim_0 by \sim_{s-1} in the above and attempting to re-run the argument. A little further thought shows this is not quite the right thing to do, since

- the space $\mathcal{E}(X)$ is too large when $s > 1$: we want a pair to represent the element of the structure group that takes one to the other, but this only makes sense for pairs $[x, y]$ that lie in the same fiber of f ;
- the condition $[x, y] \sim_{s-1} [x', y']$ turns out to be too restrictive on Y , since it implies that x and x' lie in the same fiber of f , and hence we cannot relate pairs lying over different fibers.

However, a second attempt at this does work, and generalizes Theorem 5.4 fairly cleanly.

We define

$$\mathcal{M} = \mathcal{M}_f(X) = \{[x, y] \in X \times X : f(x) = f(y)\}$$

i.e. the space of pairs that lie in the same fiber of f , and specify an equivalence relation \sim on \mathcal{M} by $[x, x'] \sim [y, y']$ if and only if $[x, y] \sim_{s-1} [x', y']$ in $\mathcal{E}(X)$, where \sim_{s-1} denotes the canonical equivalence relation on that space. Note that this definition involves a “90° rotation”, e.g. pairing up $[x, y]$ rather than $[x, x']$; this is not an error and is necessary for the definition to make sense.

Note that, by Lemma 7.16, this is equivalent to saying $[x, x'] \sim [y, y']$ if and only if $[\mathbb{L}^s(x; x'), \mathbb{L}^s(y; y')]$ is a cube of X .

We will define the structure group $A = A_f$ to be \mathcal{M}/\sim . As above, addition on A will be given by concatenation of edges, negation by $[x, y] \mapsto [y, x]$ and the action on X by a cube completion.

We now turn to the details.

Claim. *The relation \sim is a closed equivalence relation on \mathcal{M} .*

Proof of claim. For reflexivity we are asked to verify that for all $[x, x'] \in \mathcal{M}$ the configuration $[\mathbb{L}^s(x; x'), \mathbb{L}^s(x; x')]$ is an $(s+1)$ -cube of X ; equivalently by duplication, that $\mathbb{L}^s(x; x')$ is an s -cube. Since $f(x) = f(x')$ this is immediate from the assumption of relative s -ergodicity.

Symmetry is clear: if $[\mathbb{L}^s(x; x'), \mathbb{L}^s(y; y')]$ is a cube then so is $[\mathbb{L}^s(y; y'), \mathbb{L}^s(x; x')]$.

For transitivity, suppose $[x, x'], [y, y'], [z, z'] \in \mathcal{M}$ and $[x, x'] \sim [y, y']$, $[y, y'] \sim [z, z']$. Equivalently, $[\mathbb{L}^s(x; x'), \mathbb{L}^s(y; y')]$ and $[\mathbb{L}^s(y; y'), \mathbb{L}^s(z; z')]$ are cubes. It is immediate by glueing that $[\mathbb{L}^s(x; x'), \mathbb{L}^s(z; z')]$ is a cube, and hence $[x, x'] \sim [z, z']$ as required.

Once again, the topological properties are an easy check. □

Hence, as a set, $A = \mathcal{M}/\sim$ is well-defined. As in the proof of Theorem 5.4, we have considerable freedom to choose representatives for $a \in A$.

Claim. *Given $a \in A$ and $x \in X$, there exists a unique $x' \in X$ such that $a \sim [x, x']$.*

Proof of claim. Suppose $[y, y']$ is any representative of a . Necessarily $f(y) = f(y')$. Hence the configuration $[\sqcup^s(y; y'), \sqcup(x; -)]$ is an $(s+1)$ -corner of X lying above the cube $[\square^s(f(y)), \square^s(f(x))]$ of Y (which is a cube by ergodicity of Y and duplication): the fact that it is a corner is trivial by s -ergodicity of f . Since f is a fibration of degree at most s , there is a unique $x' \in f^{-1}(f(x))$ completing this $(s+1)$ -corner to a cube of X , as required. \square

Hence we may now define addition on A by concatenation of edges as before. Given $a, b \in A$ we may fix $e \in X$ and choose representatives $[x, e] \sim a$, $[e, y] \sim b$ and define $a + b$ to be the class of $[x, y]$. The identity is the class of $[e, e]$ for any e , and negation by sending $[x, y] \mapsto [y, x]$.

Claim. *These operations are well-defined, and give A the structure of a (topological) abelian group.*

Proof of claim. We first check well-definedness: that is, that these operations do not depend on the choice of representatives, or – in the case of addition – on the choice of e .

It is clear from the definition of \sim (in terms of \sim_s in $\mathcal{E}(X)$) that $[x, x] \sim [y, y]$ for any $x, y \in X$ and hence the identity class really is a class. Similarly, if $[x, x'] \sim [y, y']$ then $[x', x] \sim [y', y]$ and hence negation is well-defined.

Now suppose $[x, y] \sim [x', y']$ and $[y, z] \sim [y', z']$; we wish to show $[x, z] \sim [x', z']$ as this will show addition is well-defined. But this says $[x, x'] \sim_{s-1} [y, y']$ and $[y, y'] \sim_{s-1} [z, z']$ in $\mathcal{E}(X)$, and hence $[x, x'] \sim_{s-1} [z, z']$ in $\mathcal{E}(X)$ by transitivity and so $[x, z] \sim [x', z']$ as required.

Associativity of addition is now clear from associativity of concatenation. Similarly, that the identity and inverses behave as they should is straightforward.

We now argue commutativity. We will first need a lemma.

Lemma 7.23. *If $x, y, z, w \in X$ all lie in the same fiber of Y and $[x, y] \sim [z, w]$ then $[x, z] \sim [y, w]$.*

Proof of lemma. Note if $s = 1$ this is trivial: both are equivalent to

$$\begin{array}{ccc} z & & w \\ & \square & \\ x & & y \end{array}$$

being a cube.

For $s > 1$, we argue as follows. Choose u in the same fiber arbitrarily. By relative s -ergodicity,

$$\begin{array}{ccc} \sqcup^{s-1}(u; z) & & \sqcup^{s-1}(u; z) \\ & \square & \\ \sqcup^{s-1}(u; x) & & \sqcup^{s-1}(u; x) \end{array}$$

is a cube in X .

Note that $[x, y] \sim [z, w]$ is equivalent to $[x, z] \sim_{s-1} [y, w]$ in $\mathcal{E}(X)$. Using the universal replacement property, we can replace one of the $[x, z]$ edges by $[y, w]$ and find that $[x, y] \sim [z, w]$ is equivalent to

$$\begin{array}{ccc} \sqcup^{s-1}(u; z) & & \sqcup^{s-1}(u; w) \\ & \square & \\ \sqcup^{s-1}(u; x) & & \sqcup^{s-1}(u; y) \end{array}$$

being a cube in X . Exchanging y and z , we see that $[x, z] \sim [y, w]$ is equivalent to the same statement. \square

So, now suppose we have $a, b \in A$, and pick representatives $a \sim [x, e]$, $b \sim [e, y]$ and hence $a + b \sim [x, y]$. Further pick e' such that $[x, e'] \sim b$. Note x, y, e, e' all lie in the same fiber and $[x, e'] \sim [e, y]$, so by the lemma $[x, e] \sim [e', y]$. But then by concatenation,

$$b + a \sim [x, e'] + [e', y] \sim [x, y] \sim a + b$$

as required.

Once again, arguing continuity of these operations is straightforward from closedness of the equivalence relations, closedness of $C^k(X)$, continuity of f and compactness. This concludes the proof of the claim. \square

We now define the group action of A on X . Given $a \in A$ and $x \in X$ we take $a(x)$ (also notated $a.x$) to be the unique element of X such that $a \sim [x, a(x)]$. It is clear $0 \in A$ acts trivially. To confirm it is an action, we consider that $a \sim [x, a(x)]$ and $a' \sim [a(x), a'(a(x))]$ but then by definition of addition $a + a' \sim [x, a'(a(x))]$.

Moreover, it is clear that this action respects f (i.e. x and $a(x)$ lie in the same fiber), and that it is simply transitive on fibers of f . Again, continuity of the action is straightforward. Hence we have shown part (i).

We now turn to part (ii), which describes the cubes $C^k(X)$ in terms of $C^k(Y)$ and A . Recall that we wanted to show that the action of A on X induces a pointwise action of $C^k(\mathcal{D}_s(A))$ on $C^k(X)$ whose orbits are precisely the fibers of $f: C^k(X) \rightarrow C^k(Y)$. This can be rephrased as follows.

Claim. Suppose $k \geq 0$, $c \in C^k(X)$ and $c': \{0, 1\}^k \rightarrow X$ is another configuration such that $f(c) = f(c')$. Define a configuration $a: \{0, 1\}^k \rightarrow A$ by

$$a(\omega) \sim [c(\omega), c'(\omega)] .$$

Then c' is a cube if and only if $a \in C^{k+1}(\mathcal{D}_s(A))$.

Recall again that the cubespace $\mathcal{D}_s(A)$ was defined by the Host–Kra construction applied to the compact abelian group A with the filtration $A_0 = \cdots = A_s \supseteq A_{s+1} = \{0\}$. We note the following features:

- it is s -ergodic, i.e. every element of $A^{\{0,1\}^k}$ is a cube for $0 \leq k \leq s$; and
- it is a nilspace of degree s .

Proof of claim. We will consider three cases.

Case 1. $0 \leq k \leq s$. By relative s -ergodicity, we have that c' is always a cube of X . But by s -ergodicity of $\mathcal{D}_s(A)$, a is always a k -cube, so this is consistent.

Case 2. $k = s + 1$. For fixed c , let $T \subseteq A^{\{0,1\}^{s+1}}$ denote the set of all elements $a: \{0,1\}^{s+1} \rightarrow A$ such that $c' = a.c$ is a cube; so our goal is to show $T = C^{s+1}(\mathcal{D}_s(A))$. Then we claim

- (i) for $a \in T$, and for $F \subseteq \{0,1\}^{s+1}$ a face of dimension 1 and $b \in A$, we have $a + [b]_F \in T$; ⁸
- (ii) if $a \in T$ and $a(\omega) = 0$ for all $\omega \neq \vec{1}$, then $a = 0$.

To see (i), note that we can think of any cube c' of dimension $(s+1)$ as a cube of dimension s in the edge cubespace $\mathcal{E}(X)$. Let $[x, y]$ denote the edge corresponding to F in c' . Then $[x, b.x] \sim [y, b.y] \sim b$ in \mathcal{M} and so $[x, y] \sim_{s-1} [b.x, b.y]$ in \mathcal{E} . By the universal replacement property of \sim_{s-1} (Proposition 7.12) applied to $\mathcal{E}(X)$ (which certainly has the glueing property), replacing the vertex $[x, y]$ of the s -cube c' of $\mathcal{E}(X)$ by $[b.x, b.y]$ yields another s -cube, as required.

Part (ii) is immediate from relative $(s+1)$ -uniqueness.

Now, from (i) it follows that T is a union of cosets of $C^{s+1}(\mathcal{D}_s(A))$, since such $[b]_F$ generate $C^{s+1}(\mathcal{D}_s(A))$. In particular this includes $C^{s+1}(\mathcal{D}_s(A))$ itself as trivially $0 \in T$. Now if $a \in T$, by completing the $(s+1)$ -corner $a|_{\{0,1\} \setminus \{\vec{1}\}}$ in the nilspace $\mathcal{D}_s(A)$, we can find an element $a' \in \mathcal{D}_s(A)$ such that $a'(\omega) = a(\omega)$ for all $\omega \neq \vec{1}$. So $a - a' \in T$; but by (ii) this means $a = a'$ and hence $a \in C^{s+1}(\mathcal{D}_s(A))$.

Case 3. $k > s + 1$. Given c, c', a as in the statement, we note that c' is a cube if and only if every face of dimension $(s+1)$ is a cube (by Proposition 7.21), if and only if every face of a of dimension $(s+1)$ is a cube (by the previous case), if and only if a is a cube of $\mathcal{D}_s(A)$ (by Proposition 3.11).

□

This concludes the proof of Theorem 7.19.

□

APPENDIX A. AN EXTENDED EXPOSITION OF HOST–KRA CUBES

Here, we provide a fairly detailed exposition of the theory of the Host–Kra construction of cubes (or if you prefer, parallelepipeds) in nilpotent groups and nilmanifolds. Much of this is modelled closely on [GT10, Appendix E]; however, we provide some further results and closely worked examples.

⁸Recall (Definition 2.2) that $[b]_F$ denotes a configuration that is equal to b on F and zero elsewhere.

A.1. Some examples of filtered groups. We recall from Section 2 the definition of a filtered group. The following are some examples that will be referred to several times in what follows.

Example A.1. The simplest case is something like $G = \mathbb{R}$ with $G_0 = G_1 = \mathbb{R}$ and $G_i = \{0\}$ for $i \geq 2$. This corresponds to the lower central series filtration on the abelian group \mathbb{R} .

Example A.2. Again take $G = \mathbb{R}$ and $G_0 = G_1 = G_2 = \mathbb{R}$ and $G_i = \{0\}$ for $i \geq 3$. This is now a (somewhat trivial) degree 2 filtration on \mathbb{R} .

Example A.3. Take $G = \mathbb{R}^2$, and $G_0 = G_1 = \mathbb{R}^2$, $G_2 = \{0\} \times \mathbb{R}$ and $G_i = \{0\}$ for $i \geq 3$. This is again a degree 2 filtration on an abelian group.

Example A.4. For a non-commutative case, we consider the Heisenberg group \mathcal{H} , defined as the group of upper diagonal 3×3 matrices with ones on the diagonal, i.e.

$$\mathcal{H} = \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in \mathbb{R} \right\}.$$

We give this its lower central series filtration. Note that

$$[\mathcal{H}, \mathcal{H}] = Z(\mathcal{H}) = \left\{ \begin{pmatrix} 1 & 0 & z \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} : z \in \mathbb{R} \right\}$$

and so $\mathcal{H}_0 = \mathcal{H}_1 = \mathcal{H}$, $\mathcal{H}_2 = Z(\mathcal{H})$ as above and $\mathcal{H}_i = \{\text{id}\}$ for $i \geq 3$. So again we have a filtration of degree 2.

A.2. More on $\text{HK}^k(G_\bullet)$. The definition of $\text{HK}^k(G_\bullet)$ in Section 2 is fairly natural, but at the same time relatively inexplicit and computationally unhelpful. There, we specify $\text{HK}^k(G_\bullet)$ by giving an explicit generating set as a subgroup of $G^{\{0,1\}^k}$; but because we could in principle take arbitrarily long products in these generators, this does not give an algorithm for determining whether a given configuration $\{0,1\}^k \rightarrow G$ lies in the Host–Kra cube group or not.

It turns out that both these problems can be solved fairly straightforwardly: we can obtain a completely explicit description of $\text{HK}^k(G_\bullet)$ in terms of bounded length products, and at the same time obtain such an algorithm.

We recall some notation. If $S \subseteq [k]$, we write F_S for the face of $\{0,1\}^k$ that consists of the vertices ω with $S \subseteq \omega$. If $x \in G$ and $F \subseteq \{0,1\}^k$, we write $[x]_F$ for the element of $G^{\{0,1\}^k}$ defined by $\omega \mapsto x$ if $\omega \in F$ and $\omega \mapsto \text{id}$ otherwise.

Proposition A.5. *Take any integer $k \geq 0$, and fix an ordering $S_1 = \emptyset, \dots, S_{2^k} = \{0,1\}^k$ for the subsets of $[k]$ that respects inclusion, i.e. if $S_i \subseteq S_j$ then $i \leq j$. Then any configuration $(g_\omega)_{\omega \in \{0,1\}^k} \in G^{\{0,1\}^k}$ has a unique representation as an ordered product*

$$\prod_{i=1}^{2^k} [x_i]_{F_{S_i}} = [x_1]_{F_{S_1}} \cdots [x_{2^k}]_{F_{S_{2^k}}}$$

for parameters $x_i \in G$, $1 \leq i \leq 2^k$. The value of x_i is a function only of $\{g_\omega : \omega \subseteq S_i\}$, and moreover can be expressed as a fixed word of bounded length in these elements. Similarly g_ω depends only on $\{x_i : S_i \subseteq \omega\}$.

Finally, we have that $(g_\omega) \in \text{HK}^k(G_\bullet)$ if and only if $x_i \in G_{|S_i|}$ for all i . Hence, there is a homeomorphism

$$\begin{aligned} \text{HK}^k(G_\bullet) &\cong \prod_{i=1}^{2^k} G_{|S_i|} \\ &\cong \prod_{r=1}^k G_r^{\binom{k}{r}}. \end{aligned}$$

However, this is not an isomorphism of groups.

Proof sketch. We will give the key ideas of the proof; for any missing details, see [Gre15].

The first part is fairly straightforward. Given (g_ω) , we can solve for the variables x_i recursively in order from 1 to 2^k . Indeed, suppose $1 \leq i \leq 2^k$ and x_j have already been chosen for $j < i$. It is clear that we must have

$$g_{S_i} = \left(\prod_{j \in \{1, \dots, i-1\}, S_j \subsetneq S_i} x_j \right) \cdot x_i$$

and we can achieve this uniquely by setting

$$x_i := \left(\prod_{j \in \{1, \dots, i-1\}, S_j \subsetneq S_i} x_j \right)^{-1} g_{S_i}.$$

It is clear inductively from this construction that x_i only depends on $\{g_\omega : \omega \subseteq S_i\}$.

Now we consider the statement about $\text{HK}^k(G_\bullet)$. Certainly any element of the stated form with $x_i \in G_{|S_i|}$ is in $\text{HK}^k(G_\bullet)$; conversely, since all the generators of $\text{HK}^k(G_\bullet)$ have this form, we will be done if we can show that the set of such elements is closed under multiplication.

In particular, it would suffice to show that for any $1 \leq j \leq 2^k$ and $y \in G_{|S_j|}$, we have

$$\left(\prod_{i=1}^{2^k} [x_i]_{F_{S_i}} \right) [y]_{F_{S_j}} = \prod_{i=1}^{2^k} [x'_i]_{F_{S_i}}$$

for some new choice of variables x'_i .

We note the obvious product identity

$$[x]_F [y]_F = [xy]_F$$

and the slightly less obvious commutator identity

$$[x]_F [y]_{F'} [x]_F^{-1} [y]_{F'}^{-1} = [xyx^{-1}y^{-1}]_{F \cap F'}.$$

If $F = F_S$ and $F' = F_T$ then $F \cap F' = F_{S \cup T}$. By the filtration property, if $x \in G_{|S|}$ and $y \in G_{|T|}$ then $xyx^{-1}y^{-1} \in G_{|S|+|T|} \subseteq G_{|S \cup T|}$ and so the left hand side is again a generator of the Host–Kra cube group. Similarly if $x, y \in G_{|S|}$ then so is xy , and so $[xy]_{F_S}$ is again a generator.

This means that we can move the new term $[y]_{F_{S_j}}$ to the left in the product, generating new error terms as we go, but ones of strictly “lower order” as measured by the larger size of the set S concerned. Arguing by induction on $|S_j|$, we can iteratively clean up these errors, and the process clearly terminates when $|S_j| = k$. \square

Remark A.6. One can be completely explicit about the various expressions that arise in this proof, e.g. the expressions for x_i in terms of (g_ω) . These turn out to be products with alternating inverse signs, traversing “Gray codes” around the vertices of $\{0, 1\}^k$. This observation was first made in [ACS12, Section 1.2]; the reader could also consult [Gre15] for an exposition.

This Proposition allows us to compute some spaces of Host–Kra cubes in the case of our example filtered groups.

Example A.7. Consider Example A.1. Then $\mathrm{HK}^0(G) = G$, $\mathrm{HK}^1(G) = G^2$ and

$$\mathrm{HK}^2(G) = \{(x, y, z, w) \in G^4 : x - y - z + w = 0\}$$

i.e. the Host–Kra 2-cubes correspond exactly to parallelograms. Similarly, $\mathrm{HK}^k(G)$ consists precisely of parallelepipeds of dimension k in the usual sense.

Example A.8. In Example A.2, we have $\mathrm{HK}^0(G) = G$, $\mathrm{HK}^1(G) = G^2$ and also $\mathrm{HK}^2(G) = G^4$. The interesting case is $\mathrm{HK}^3(G)$ which is given by

$$\mathrm{HK}^3(G) = \left\{ (x_{000}, \dots, x_{111}) \in \mathbb{R}^{\{0,1\}^3} : x_{000} - x_{001} - x_{010} + x_{011} - \dots - x_{111} = 0 \right\}$$

i.e. all tuples whose alternating sum is zero. Indeed, if we write (x_ω) as

$$\sum_{i=1}^8 [y_i]_{F_{S_i}}$$

as in the proposition, the conditions on y_i just reduce to $y_8 = 0$. It is not hard to extract the explicit formula (in the abelian case)

$$y_8 = x_{000} - x_{001} - x_{010} + x_{011} - \dots - x_{111}$$

from the proof of the proposition.

Example A.9. We consider the degree s filtration on an Abelian group

$$A = A_0 = A_1 = \dots = A_s \supseteq A_{s+1} = \{0\},$$

because of the important role it plays in the weak structure theory.

Specifically, we verify the claim made in Proposition 5.1 that a configuration $c : \{0, 1\}^{s+1} \rightarrow A$ is a cube in $\mathcal{D}_s(A)$ (i.e. $c \in \mathrm{HK}^{s+1}(A_\bullet)$) if and only if ⁹

$$L(c) := \sum_{\omega \in \{0,1\}^{s+1}} (-1)^{|\omega|} c(\omega) = 0.$$

⁹Here $|\omega|$ denotes the number of 1 entries in ω .

We first observe that $L(c)$ is an additive function on configurations, since the group is commutative. If we write $c = \sum [x_i]_{F_{S_i}}$ as in Proposition A.5, we get

$$L(c) = \sum L([x_i]_{F_{S_i}}).$$

We note that $L([x]_F) = 0$ for any face F of dimension at least 1 and $L([x]_F) = \pm x$ for a face F of dimension 0, as follows directly from the definition of L . Thus $L(c) = \pm x_{2^n}$. Now the claim follows from Proposition A.5, which asserts that $c \in \text{HK}^s(A_\bullet)$ if and only if $x_{2^n} \in A_{s+1} = \{0\}$.

Example A.10. For Example A.3, the Host–Kra k -cubes are the direct product of the k -cubes from Examples A.7 and A.8. In particular, $\text{HK}^3(G)$ consists of tuples (x_ω, y_ω) such that (x_ω) form a parallelepiped and the alternating sum of (y_ω) is zero.

Example A.11. For the Heisenberg group (Example A.4), despite the non-abelian group law the situation is similar to the previous example. Certainly $\text{HK}^0 = \mathcal{H}$, $\text{HK}^1 = \mathcal{H}^2$ as usual. A configuration

$$\omega \mapsto \begin{pmatrix} 1 & x_\omega & z_\omega \\ 0 & 1 & y_\omega \\ 0 & 0 & 1 \end{pmatrix}$$

is in $\text{HK}^2(\mathcal{H}_\bullet)$ if and only if it has the form

$$[g_1]_{F_\emptyset} [g_2]_{F_{\{1\}}} [g_3]_{F_{\{2\}}} [g_4]_{F_{\{1,2\}}}$$

for some $g_1, g_2, g_3 \in \mathcal{H}$ and $g_4 \in \mathcal{H}_2$. In the “abelian” variables x_ω, y_ω this condition becomes $x_{00} - x_{01} - x_{10} + x_{11} = 0$ and $y_{00} - y_{01} - y_{10} + y_{11} = 0$; and for z_ω we have complete freedom.

Now we consider

$$\omega \mapsto \begin{pmatrix} 1 & x_\omega & z_\omega \\ 0 & 1 & y_\omega \\ 0 & 0 & 1 \end{pmatrix}$$

which it turns out is in $\text{HK}^3(\mathcal{H}_\bullet)$ if and only if:

- (x_ω) is a parallelepiped;
- (y_ω) is a parallelepiped;
- the alternating sum $\sum_{\omega \in \{0,1\}^3} (-1)^{|\omega|} z_\omega$ is zero.

The first two points are fairly easy to justify: since this calculation really takes place on the abelianization of \mathcal{H} which is just \mathbb{R}^2 , this strongly resembles Example A.7.

For the third, we again use a decomposition

$$\prod_{i=1}^8 [g_i]_{F_{S_i}}$$

and consider the constraint $g_8 = \text{id}$ and in particular its z -coordinate. As in Example A.8, the contribution from z_ω is just this alternating sum. It is not immediately obvious that this expression should not also involve terms of the form $x_\omega y_\nu$ arising from the cross-terms in the group operation, but a calculation shows that these do indeed drop out.

A.3. Some topological properties of the Host–Kra groups. In the second paper of this project [GMV16a], we will use some topological properties of the Host–Kra groups. These are easy to deduce from Proposition A.5, so we record them now for convenient reference.

Lemma A.12. *Let G_\bullet, H_\bullet be two filtered topological groups. Let $\tau : G \rightarrow H$ be a homomorphism such that $\tau(G_i) \subseteq \tau(H_i)$. Then τ induces a homomorphism $\tau : \mathrm{HK}^n(G_\bullet) \rightarrow \mathrm{HK}^n(H_\bullet)$ for each n by pointwise application on the vertices.*

If $\tau : G_i \rightarrow H_i$ is open for each i , then so is the induced homomorphism $\tau : \mathrm{HK}^n(G_\bullet) \rightarrow \mathrm{HK}^n(H_\bullet)$ for each n .

Lemma A.13. *Let G_\bullet, H_\bullet be two filtered topological groups. Suppose $G_i \subseteq H_i$ for each i . If G_i are open in H_i (resp., connected) for each i , then $\mathrm{HK}^n(G_\bullet)$ is also open in $\mathrm{HK}^n(H_\bullet)$ (resp., connected) for each n .*

We recall from Proposition A.5 that each element (g_ω) can be written uniquely in the form $\prod [x_i]_{F_{S_i}}$. For convenient reference, we call (g_ω) the *vertex coordinates* and (x_i) the *face coordinates*. We also recall that the two set of coordinates can be expressed from each other by word maps of finite length.

Proof of Lemma A.12. That $\tau(\mathrm{HK}^n(G_\bullet)) \subseteq \mathrm{HK}^n(H_\bullet)$ is clear by considering the generators.

We observe that the face coordinates give rise to the same homomorphism by pointwise application, since the two sets of coordinates can be expressed from each other using word maps. The last claim follows from the description with face coordinates, since a product of open maps is open. \square

Proof of Lemma A.13. We use the description with face coordinates and the fact that the direct product of open (resp., connected) sets is open (resp., connected). \square

A.4. The Host–Kra group $\mathrm{HK}^k(G_\bullet)$ is a nilspace. Another crucial corollary of Proposition A.5 is a version of the *corner constraint*.

Corollary A.14. *Suppose G_\bullet is a degree s filtered group, set $k = s + 1$ and suppose $g, g' \in \mathrm{HK}^k(G_\bullet)$ have the property that $g(\omega) = g'(\omega)$ for all $\omega \neq \vec{1}$.¹⁰*

Then $g(\vec{1}) = g'(\vec{1})$.

Proof. Apply Proposition A.5 to g and g' to obtain parameters x_i, x'_i . Every x_i other than the last one x_{2^k} does not depend on $g(\vec{1})$, and similarly for g' . But $G_k = \{\mathrm{id}\}$ by assumption, so $x_{2^k} = x'_{2^k} = \mathrm{id}$ in the decomposition. Hence $g = g'$. \square

We consider the *corner completion* property discussed at length above.

Corollary A.15. *Suppose we have a configuration $(g_\omega)_{\omega \in \{0,1\}^k \setminus \{\vec{1}\}}$ such that every lower face is in $\mathrm{HK}^{k-1}(G_\bullet)$. Then there exists $g \in G$ such that setting $g_{\vec{1}} = g$ we have $(g_\omega)_{\omega \in \{0,1\}^k} \in \mathrm{HK}^k(G_\bullet)$. In other words, the cubespace $(G, \mathrm{HK}^k(G_\bullet))$ is fibrant.*

¹⁰As ever, we write $\vec{1}$ for the element $(1, \dots, 1) \in \{0, 1\}^k$.

Note the previous corollary says precisely that the choice of g is unique when $k = s + 1$.

Proof. By Proposition A.5 we may define all the variables x_i for $1 \leq i < 2^k$ without knowing $g_{\vec{1}}$, such that

$$(g_\omega) = \prod_{i=1}^{2^k-1} [x_i]_{F_{S_i}}$$

whenever this is defined.

We show that $x_j \in G_{|S_j|}$. To this end, we choose a lower face E that contains the vertex S_j and consider the product

$$\prod_{i \in \{1, \dots, 2^k-1\}: S_i \in E}^{2^k-1} [x_i]_{F_{S_i}}|_E,$$

where $\cdot|_E$ denotes restriction to E . We observe that this is a product decomposition of a cube in $\mathrm{HK}^{\dim E}(G_\bullet)$ of the form that appears in Proposition A.5. The second part of that proposition now implies that $x_j \in G_{|S_j|}$.

If we set $x_{2^k} = \mathrm{id}$, then

$$\prod_{i=1}^{2^k-1} [x_i]_{F_{S_i}}$$

is an element of $\mathrm{HK}^k(G_\bullet)$ extending (g_ω) as required. \square

A.5. Properties of cubes on nilmanifolds. We recall from Definition 2.4 that the notion of Host–Kra cubes on G_\bullet induces one on nilmanifolds G/Γ .

Many of the nice properties of the Host–Kra cubes on G are inherited on the nilmanifold G/Γ , even though the latter typically has no compatible group operation. We will verify some instances of this that are of particular importance in the abstract setting. We begin by showing that the Host–Kra cubes arising from a filtration of degree s satisfy $(s+1)$ -uniqueness.

Proposition A.16. *Take G_\bullet, Γ as in the definition, where the filtration G_\bullet has degree s . Set $k = s+1$, and suppose $c, c' \in \mathrm{HK}^k(G_\bullet)/\Gamma$ have the property that $c(\omega) = c'(\omega)$ for all $\omega \neq \vec{1}$. Then $c = c'$.*

Proof. We write $\pi: G \rightarrow G/\Gamma$ for the projection map. Let $g, g' \in \mathrm{HK}^k(G_\bullet)$ be such that $\pi(g) = c$, $\pi(g') = c'$. Hence $g^{-1}g'$ is in $\mathrm{HK}^k(G_\bullet)$, and all of its entries lie in Γ except possibly the $\vec{1}$ entry. Applying Proposition A.5, we find that all the coefficients x_i in the expansion

$$g^{-1}g' = \prod_{i=1}^{2^k} [x_i]_{F_{S_i}}$$

lie in Γ , except possibly the last one x_{2^k} ; but in fact $x_{2^k} \in G_k = \{\mathrm{id}\}$ and so x_{2^k} is in Γ also. Hence, $(g^{-1}g')(\vec{1}) \in \Gamma$, and so $g(\vec{1})\Gamma = g'(\vec{1})\Gamma$ as required. \square

We now turn to the completion property.

Proposition A.17. *Take G_\bullet, Γ as above, and suppose $(y_\omega)_{\omega \in \{0,1\}^k \setminus \{\vec{1}\}}$ is a configuration on G/Γ such that every lower face is in $\mathrm{HK}^{k-1}(G_\bullet)/\Gamma$. Then there exists $y \in G/\Gamma$ such that setting $y_{\vec{1}} = y$ results in an element of $\mathrm{HK}^k(G_\bullet)/\Gamma$.*

Proof. Since $(G, \mathrm{HK}^k(G_\bullet))$ is fibrant (Corollary A.15), it suffices by the universal property (Lemma 7.8) to verify that the projection $G \rightarrow G/\Gamma$ is a fibration.

Unpacking the definitions, this says the following. Suppose we are given a cube $c \in \mathrm{HK}^k(G_\bullet)/\Gamma$; by definition this has the form $\omega \mapsto g(\omega)\Gamma$ where $\omega \mapsto g(\omega)$ is a Host–Kra cube of G_\bullet . Suppose we are given another k -corner

$$h: \{0,1\}^k \setminus \{\vec{1}\} \rightarrow G$$

such that $h(\omega)\Gamma = g(\omega)\Gamma$ for all $\omega \neq \vec{1}$. Then we need to show there is an $h(\vec{1}) \in G$ such that $h(\vec{1})\Gamma = g(\vec{1})\Gamma$, and furthermore this completes h to a cube of G .

Consider the k -corner $b: \omega \mapsto g(\omega)^{-1}h(\omega)$. By assumption this takes values in Γ . By Proposition A.5 we may write b as an ordered product

$$\prod_{i=1}^{2^k} [x_i]_{F_{S_i}}$$

for some $x_i \in G_{|S_i|}$, where all but x_{2^k} are determined. Moreover, all of these elements lie in Γ . Setting $x_{2^k} = \mathrm{id}$ (say) we complete b to a cube in $\mathrm{HK}^k(G_\bullet)$ with entries in Γ , and then we can set $h(\vec{1}) = g(\vec{1})b(\vec{1})$ as required. \square

A.6. Examples of cubes on nilmanifolds. Our next task is to explore the consequences of Definition 2.4 concerning cubes on nilmanifolds, in the setting of our favourite examples.

The abelian cases here are all fairly straightforward.

Example A.18. Suppose we take $\Gamma = \mathbb{Z}$ in Examples A.1 or A.2, or $\Gamma = \mathbb{Z}^2$ in Example A.3. Then G/Γ actually does have a group structure in these cases; specifically, it is the torus \mathbb{R}/\mathbb{Z} or $\mathbb{R}^2/\mathbb{Z}^2$ respectively. Moreover, this compact group inherits a filtration from that on \mathbb{R} or \mathbb{R}^2 .

In this case, the Host–Kra cubes on the “nilmanifold” G/Γ agree with the Host–Kra construction on the group G/Γ .

In the case of the Heisenberg group (Example A.4) and the corresponding nilmanifolds, it is much harder to obtain a completely explicit description of the Host–Kra cubespaces $\mathrm{HK}^k(G_\bullet)/\Gamma$. We do so anyway, both as an instructive exercise in its own right, and to emphasize some of the difficulties that arise in trying to obtain a simple description of nilspaces in general.

Example A.19. Consider again the Heisenberg case (Example A.4) with Γ the subgroup of \mathcal{H} consisting of matrices with integer entries. A fundamental domain for the quotient \mathcal{H}/Γ consists of all matrices

$$\left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in [0, 1) \right\}$$

and the reduction map to the fundamental domain is

$$\begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \mapsto \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 - \lfloor x \rfloor & -\lfloor z - x \lfloor y \rfloor \rfloor \\ 0 & 1 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & \{x\} & \{z - x \lfloor y \rfloor\} \\ 0 & 1 & \{y\} \\ 0 & 0 & 1 \end{pmatrix}$$

where $\lfloor \cdot \rfloor$, $\{ \cdot \}$ denote the integer and fractional parts respectively.

Let

$$c(\omega) = \begin{pmatrix} 1 & x_\omega & z_\omega \\ 0 & 1 & y_\omega \\ 0 & 0 & 1 \end{pmatrix}$$

be the representatives in the fundamental domain of an element of $(G/\Gamma)^{\{0,1\}^3}$. We wish to find necessary and sufficient conditions for this to represent an element of $\text{HK}^3(\mathcal{H}_\bullet)/\Gamma$.

By definition, this holds if and only if there is a configuration

$$c'(\omega) = \begin{pmatrix} 1 & r_\omega & t_\omega \\ 0 & 1 & s_\omega \\ 0 & 0 & 1 \end{pmatrix}$$

in $\text{HK}^3(\mathcal{H}_\bullet)$ such that $c(\omega)$ are precisely the representatives of $c'(\omega)$ in the fundamental domain; i.e. if

$$\begin{pmatrix} 1 & x_\omega & z_\omega \\ 0 & 1 & y_\omega \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & \{r_\omega\} & \{t_\omega - r_\omega \lfloor s_\omega \rfloor\} \\ 0 & 1 & \{s_\omega\} \\ 0 & 0 & 1 \end{pmatrix}$$

for all ω .

By the description of $\text{HK}^3(\mathcal{H}_\bullet)$ (see Example A.11) this implies that $(x_\omega \bmod 1)_{\omega \in \{0,1\}^3}$ and $(y_\omega \bmod 1)_{\omega \in \{0,1\}^3}$ must be three-dimensional parallelepipeds over \mathbb{R}/\mathbb{Z} . One might expect a corresponding identity

$$z_{000} - z_{001} - z_{010} + z_{011} - \cdots - z_{111} \equiv 0 \pmod{1}$$

as was the case for $\text{HK}^3(\mathcal{H}_\bullet)$; but in fact this gets “twisted” by the non-abelian group action. For instance, taking the cube

$$c(\omega) = \begin{pmatrix} 1 & 1/2 & 0 \\ 0 & 1 & (\omega_1 + \omega_2 + \omega_3)/3 \\ 0 & 0 & 1 \end{pmatrix}$$

in $\text{HK}^3(\mathcal{H}_\bullet)$, its projection to the fundamental domain is

$$c'(\omega) = \begin{cases} \begin{pmatrix} 1 & 1/2 & 0 \\ 0 & 1 & (\omega_1 + \omega_2 + \omega_3)/3 \\ 0 & 0 & 1 \end{pmatrix} & : \omega \neq \vec{1} \\ \begin{pmatrix} 1 & 1/2 & 1/2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} & : \omega = \vec{1} \end{cases}$$

and so the alternating sum over the z coordinate is $1/2 \bmod 1$. However, once this value $1/2$ as a function of x_ω, y_ω is fixed, it is true that e.g.

$$c''(\omega) = \begin{pmatrix} 1 & 1/2 & z_\omega \\ 0 & 1 & 1/3(\omega_1 + \omega_2 + \omega_3) \bmod 1 \\ 0 & 0 & 1 \end{pmatrix}$$

is in $\text{HK}^3(\mathcal{H}_\bullet)/\Gamma$ if and only if $z_{000} - z_{001} - \cdots + z_{110} - z_{111} = 1/2 \pmod{1}$.

It is possible but not obviously informative to do the calculation in general. We compute

$$\begin{aligned} \sum_{\omega} (-1)^{|\omega|} z_{\omega} &\equiv \sum_{\omega} (-1)^{|\omega|} (t_{\omega} - r_{\omega} \lfloor s_{\omega} \rfloor) \pmod{1} \\ &\equiv - \sum_{\omega} (-1)^{|\omega|} r_{\omega} \lfloor s_{\omega} \rfloor \pmod{1} \end{aligned}$$

since the alternating sum over t_{ω} is known to be zero. Since $r_{\omega} - x_{\omega}$ is an integer, we can rewrite this as

$$- \sum_{\omega} (-1)^{|\omega|} x_{\omega} \lfloor s_{\omega} \rfloor \pmod{1}$$

and by exploiting the fact that $s_{\omega} \equiv y_{\omega} \pmod{1}$ together with the parallelepiped conditions for (s_{ω}) and $(x_{\omega} \pmod{1})$ and a little calculation, we can derive a totally unenlightening formula

$$\begin{aligned} \sum_{\omega} (-1)^{|\omega|} z_{\omega} &\equiv x_{110}(y_{110} - y_{100} - y_{010} + y_{000}) \\ &\quad + x_{101}(y_{101} - y_{100} - y_{001} + y_{000}) \\ &\quad + x_{011}(y_{011} - y_{010} - y_{001} + y_{000}) \\ &\quad + x_{111}(y_{111} + 2y_{000} - y_{001} - y_{010} - y_{100}) . \end{aligned}$$

Together with the previous conditions, this gives necessary and sufficient conditions for a configuration $(x_{\omega}, y_{\omega}, z_{\omega})$ to be a Host–Kra cube of \mathcal{H}/Γ .

The specific form of this explicit “cocycle” on the right hand side is not terribly important. What is useful to bear in mind is that these conditions have the structure:

- the “abelianization” $(x_{\omega} \pmod{1}, y_{\omega} \pmod{1}) \in \mathbb{R}^2/\mathbb{Z}^2$ must be a parallelepiped; and
- the alternating sum $\sum_{\omega} (-1)^{|\omega|} z_{\omega} \pmod{1}$ is given by some function only of the x ’s and y ’s.

A.7. Polynomial sequences. We finally give a brief discussion of *polynomial sequences*.

One of the key features of filtered groups is that they admit a natural definition of a *polynomial sequence*. These are a class of functions $\mathbb{Z} \rightarrow G$, generalizing the notion of a polynomial $\mathbb{Z} \rightarrow \mathbb{R}$.

Definition A.20. For any $f: \mathbb{Z} \rightarrow G$ and $h, x \in \mathbb{Z}$ write $\partial_h f: x \mapsto f(x+h)f(x)^{-1}$. A map $f: \mathbb{Z} \rightarrow G$ is called a *polynomial sequence* if $\partial_{h_i} \dots \partial_{h_1} f(x) \in G_i$ for each $i \geq 0$ and $h_1, \dots, h_i \in \mathbb{Z}$.

To justify the term “polynomial”, we consider this definition in the context of our examples.

Example A.21. In Example A.1, the polynomial sequences are precisely the polynomials of degree at most 1 on \mathbb{Z} , i.e. maps of the form $n \mapsto an + b$ for some $a, b \in \mathbb{R}$. Indeed, the condition translates to the functional equation $p(x) - p(x+h_1) - p(x+h_2) + p(x+h_1+h_2) = 0$ for all $x, h_1, h_2 \in \mathbb{Z}$, which characterizes affine-linear functions.

Example A.22. Similarly, in Example A.2, we get precisely the polynomials of degree at most 2 on \mathbb{Z} , i.e. maps of the form $n \mapsto an^2 + bn + c$ for some $a, b, c \in \mathbb{R}$. As before, the definition expands to a

functional equation for p ,

$$p(x) - p(x+h_1) - p(x+h_2) + p(x+h_1+h_2) - p(x+h_3) + p(x+h_1+h_3) + p(x+h_2+h_3) - p(x+h_1+h_2+h_3) = 0$$

which has the required properties. Alternatively, the definition holds if and only if $\partial_h p$ is a polynomial of degree 1 for all $h \in \mathbb{Z}$.

Example A.23. Example A.3 is a hybrid of these: the polynomial sequences are of the form $n \mapsto (p_1(n), p_2(n))$ where p_1 has degree at most 1 and p_2 has degree at most 2.

Example A.24. The polynomial sequences for Example A.4 have the form

$$n \mapsto \begin{pmatrix} 1 & a_1 n + a_0 & c_2 n^2 + c_1 n + c_0 \\ 0 & 1 & b_1 n + b_0 \\ 0 & 0 & 1 \end{pmatrix}$$

for any choice of $a_i, b_i, c_i \in \mathbb{R}$. This is no longer quite so clear, due to the non-abelian group operation implicit in the definition of a polynomial sequence. However, it follows from a short computation, or from Proposition A.25 below.

These sequences are fairly hard to study directly from this definition. For instance, it is a true fact that if $p(n)$ and $p'(n)$ are polynomial sequences then so is $p(n)p'(n)$, but a direct argument is fairly involved.

The theory becomes much cleaner, however, if we invoke the notion of Host–Kra cube groups defined in Definition 2.2. The key result is the following.

Proposition A.25. *Let G_\bullet be a filtered group and $p: \mathbb{Z} \rightarrow G$ be given. Then p is a polynomial sequence if and only if the following holds: for any $k \geq 0$ and $x, h_1, \dots, h_k \in \mathbb{Z}$, the configuration*

$$\begin{aligned} \{0, 1\}^k &\rightarrow G \\ \omega &\mapsto p \left(x + \sum_{i=1}^k h_i \omega_i \right) \end{aligned}$$

lies in $\text{HK}^k(G_\bullet)$.

Equivalently, p is a polynomial sequence if and only if it maps $\text{HK}^k(\mathbb{Z}_\bullet)$ into $\text{HK}^k(G_\bullet)$ for each k , where \mathbb{Z} is given the usual filtration $\mathbb{Z}_0 = \mathbb{Z}_1 = \mathbb{Z}$, $\mathbb{Z}_i = \{0\}$ for $i \geq 2$.

For a proof, see [GT12, Section 6]. Hence, the notion of a polynomial map $\mathbb{Z} \rightarrow G_\bullet$ (which is the common term in the literature) coincides with the concept from the introduction of a map that sends cubes to cubes; or with the formal notion of a cubespace morphism.

In the case of maps $\mathbb{Z} \rightarrow G/\Gamma$, no criterion along the lines of A.20 is available. The most common convention in the literature is to refer to a map $p: \mathbb{Z} \rightarrow G/\Gamma$ as polynomial if it is the projection of a polynomial map $p: \mathbb{Z} \rightarrow G$ in the existing sense.

Alternatively, one can simply use the notion of a cubespace morphism $\mathbb{Z} \rightarrow G/\Gamma$; i.e., a map that sends Host–Kra cubes of \mathbb{Z} (i.e. parallelepipeds) to Host–Kra cubes of G/Γ .

It turns out that these notions coincide. However, this is not the case if \mathbb{Z} is replaced by another group; e.g. in the case of maps $\mathbb{Z}/N\mathbb{Z} \rightarrow G/\Gamma$, we find that there are no non-constant polynomial maps $\mathbb{Z}/N\mathbb{Z} \rightarrow G$ but plenty of maps $\mathbb{Z}/N\mathbb{Z} \rightarrow G/\Gamma$ that send cubes to cubes. The reader is encouraged always to have in mind the latter notion.

REFERENCES

- [ACS12] O. Antolín Camarena and B. Szegedy, *Nilspaces, nilmanifolds and their morphisms*, 2012. arXiv:1009.3825v3. [↑2](#), [5](#), [6](#), [11](#), [15](#), [17](#), [18](#), [21](#), [22](#), [25](#), [26](#), [27](#), [30](#), [37](#), [42](#), [53](#)
- [Can16a] P. Candela, *Notes on compact nilspaces*, 2016. Preprint. [↑2](#)
- [Can16b] P. Candela, *Notes on nilspaces – algebraic aspects*, 2016. Preprint. [↑2](#)
- [GMV16a] Y. Gutman, F. Manners, and P. P. Varjú, *The structure theory of nilspaces II: Representation as nilmanifolds*, 2016. [↑1](#), [2](#), [6](#), [10](#), [18](#), [19](#), [22](#), [25](#), [55](#)
- [GMV16b] Y. Gutman, F. Manners, and P. P. Varjú, *The structure theory of nilspaces III: Inverse limit representations and topological dynamics*, 2016. [↑1](#), [2](#), [6](#), [10](#), [17](#), [19](#), [22](#), [25](#), [27](#)
- [Gre15] B. Green, *Higher-order fourier analysis I*, Notes available from the author, 2015. [↑3](#), [52](#), [53](#)
- [GT08] B. Green and T. Tao, *An inverse theorem for the Gowers $U^3(G)$ norm*, Proc. Edinb. Math. Soc. (2) **51** (2008), no. 1, 73–153. MR2391635 [↑4](#)
- [GT10] B. Green and T. Tao, *Linear equations in primes*, Ann. of Math. (2) **171** (2010), no. 3, 1753–1850. MR2680398 (2011j:11177) [↑2](#), [3](#), [7](#), [11](#), [50](#)
- [GT12] B. Green and T. Tao, *The quantitative behaviour of polynomial orbits on nilmanifolds*, Ann. of Math. (2) **175** (2012), no. 2, 465–540. MR2877065 [↑60](#)
- [HK05] B. Host and B. Kra, *Nonconventional ergodic averages and nilmanifolds*, Ann. of Math. (2) **161** (2005), no. 1, 397–488. MR2150389 (2007b:37004) [↑2](#), [3](#), [7](#)
- [HK08] B. Host and B. Kra, *Parallelepipedes, nilpotent groups and Gowers norms*, Bull. Soc. Math. France **136** (2008), no. 3, 405–437. MR2415348 (2009d:11021) [↑2](#), [3](#), [6](#), [7](#), [11](#), [21](#), [30](#)
- [HKM10] B. Host, B. Kra, and A. Maass, *Nilsequences and a structure theorem for topological dynamical systems*, Adv. Math. **224** (2010), no. 1, 103–129. MR2600993 [↑2](#)
- [HM07] B. Host and A. Maass, *Nilsystèmes d’ordre 2 et parallélépipèdes*, Bull. Soc. Math. France **135** (2007), no. 3, 367–405. MR2430186 (2010j:37013) [↑2](#)
- [MKS66] W. Magnus, A. Karrass, and D. Solitar, *Combinatorial group theory: Presentations of groups in terms of generators and relations*, Interscience Publishers [John Wiley & Sons, Inc.], New York-London-Sydney, 1966. MR0207802 (34 #7617) [↑7](#)
- [Rud95] D. J. Rudolph, *Eigenfunctions of $T \times S$ and the Conze-Lesigne algebra.*, Ergodic theory and its connections with harmonic analysis. Proceedings of the 1993 Alexandria conference, Alexandria, Egypt, May 24–28, 1993, 1995, pp. 369–432 (English). [↑18](#)
- [Sze10] B. Szegedy, *Structure of finite nilspaces and inverse theorems for the gowers norms in bounded exponent groups*, 2010. arXiv:1011.1057. [↑18](#)
- [Sze12] B. Szegedy, *On higher order fourier analysis*, 2012. arXiv:1203.2260v1. [↑2](#), [5](#)

YONATAN GUTMAN, INSTITUTE OF MATHEMATICS, POLISH ACADEMY OF SCIENCES, UL. ŚNIADECKICH 8, 00-656 WARSZAWA, POLAND.

E-mail address: y.gutman@impan.pl

FREDDIE MANNERS, MATHEMATICAL INSTITUTE, RADCLIFFE OBSERVATORY QUARTER, WOODSTOCK ROAD, OXFORD
OX2 6GG

E-mail address: `Frederick.Manners@maths.ox.ac.uk`

PÉTER P. VARJÚ, CENTRE FOR MATHEMATICAL SCIENCES, WILBERFORCE ROAD, CAMBRIDGE CB3 0WA, UK

E-mail address: `pv270@dpms.cam.ac.uk`